Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

# NEMO: Towards exa-scale ocean simulation

Oriol Tintó Prims

Advisors:

- Anna Cortés
- Mario C. Acosta
- Francisco J. Doblas-Reyes

20/04/2018

Universitat Autònoma de Barcelona

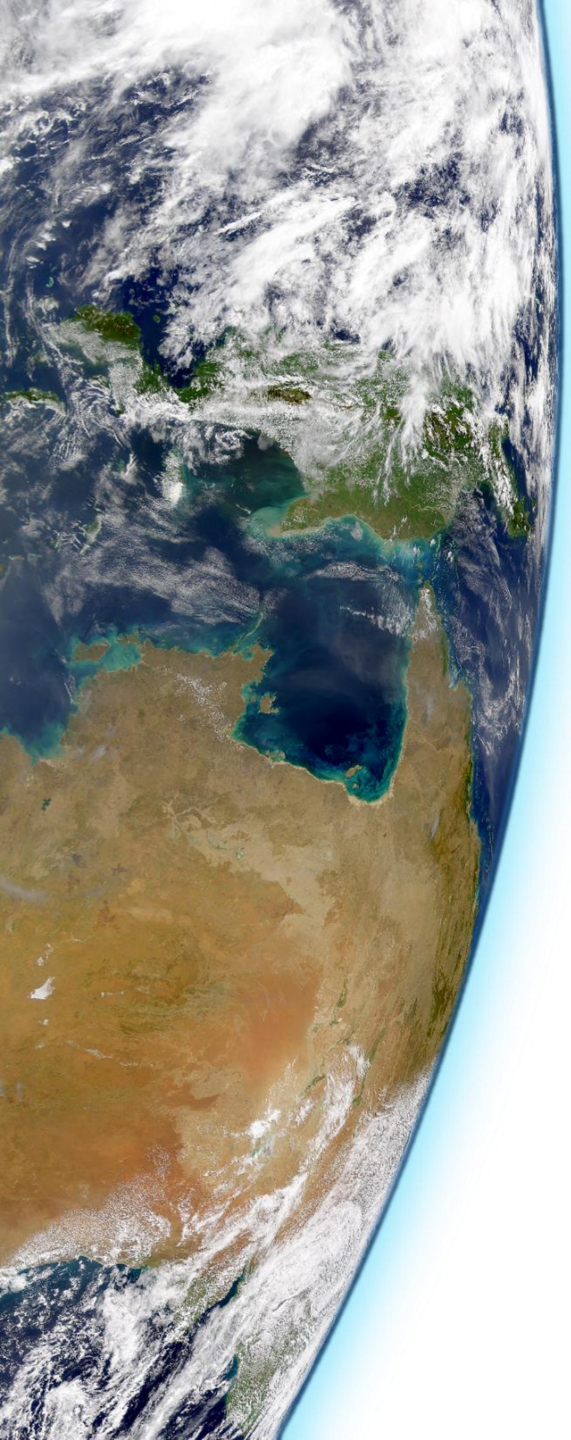

# Outline

- Introduction
- Challenges
- Work done
- Collaborations
- Conclusions

# Introduction

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

# Motivation

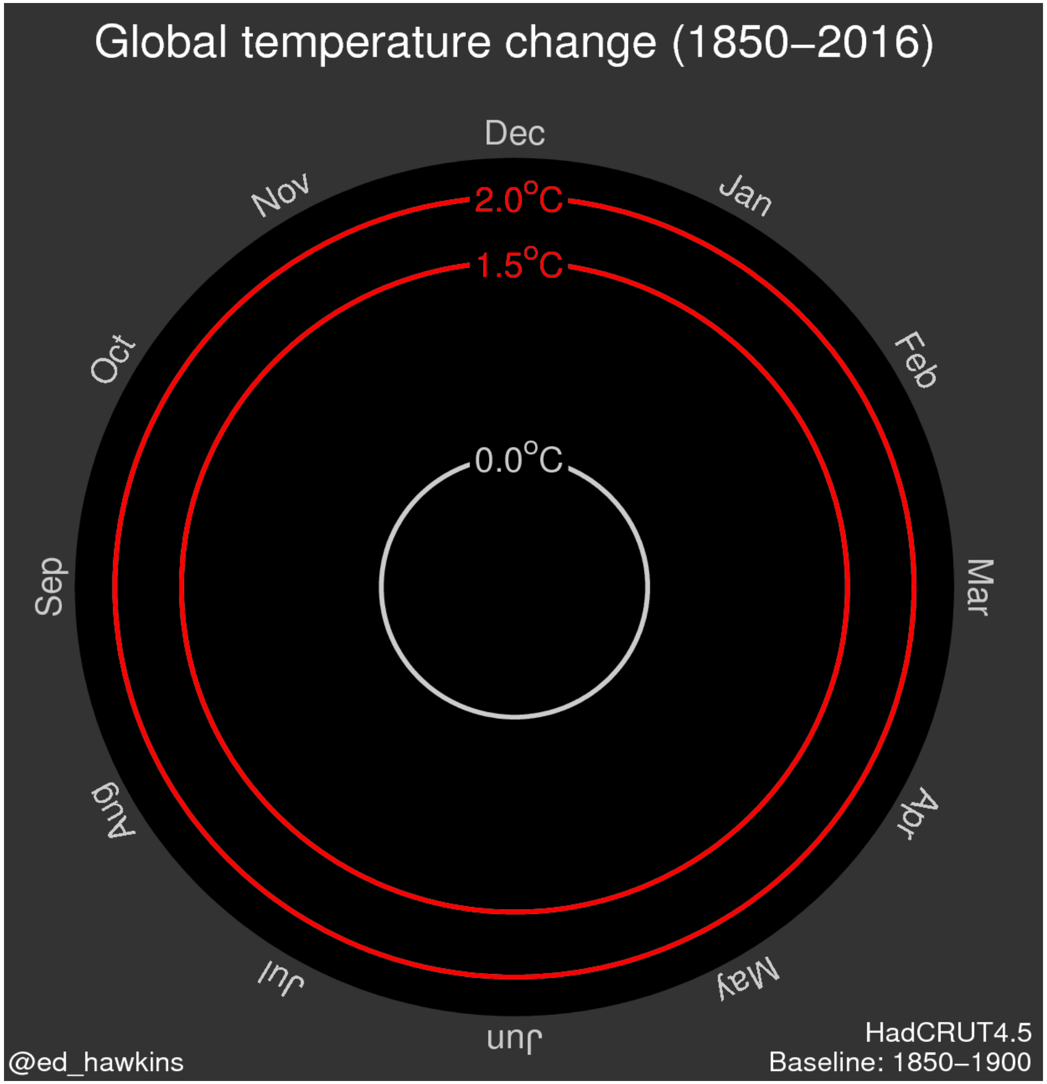

Global temperature change (1850–2016)
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
2.0°C
1.5°C
0.0°C
@ed_hawkins
HadCRUT4.5
Baseline: 1850–1900

# Motivation

# Motivation

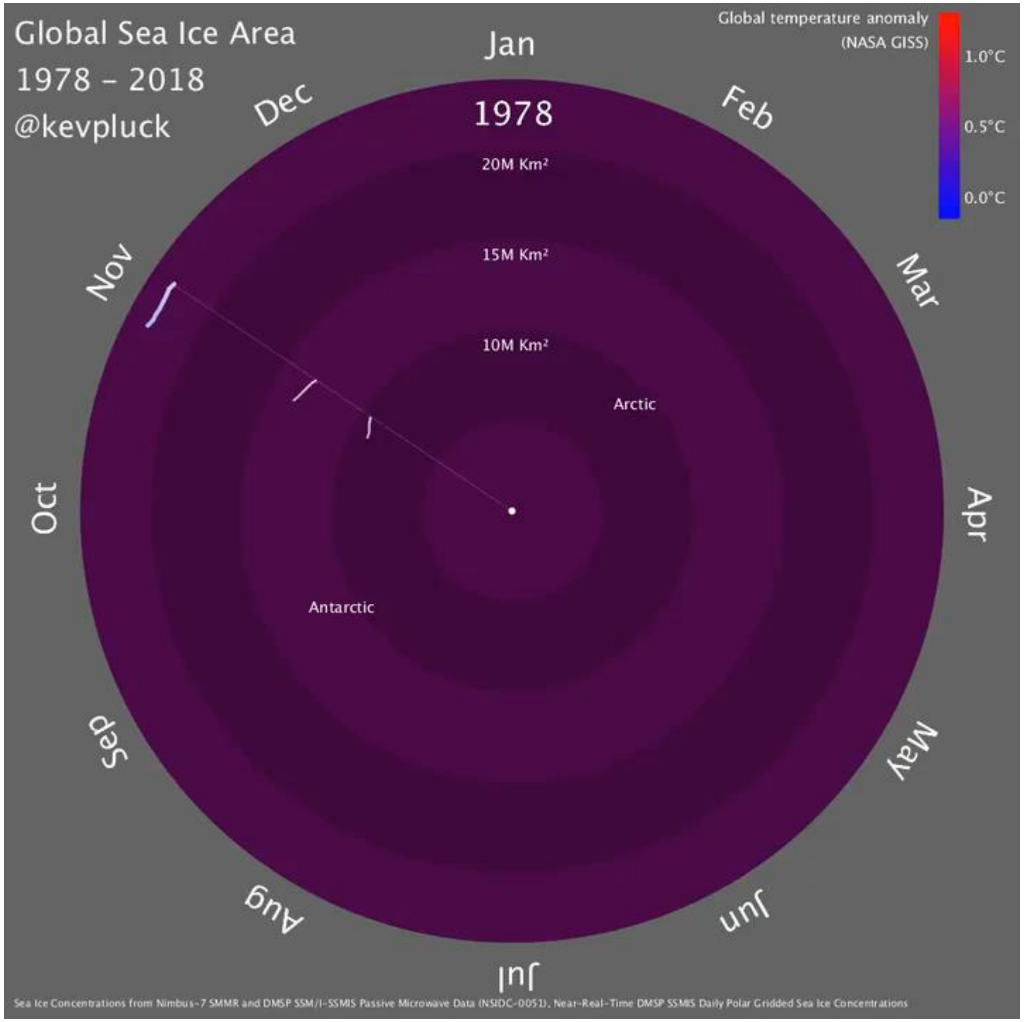


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Motivation

Neumayer Glacier shrinks on South Georgia Island
**January 11, 2005** - September 14, 2016

# Motivation

Neumayer Glacier shrinks on South Georgia Island
January 11, 2005 - **September 14, 2016**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB
Universitat Autònoma de Barcelona

# Motivation

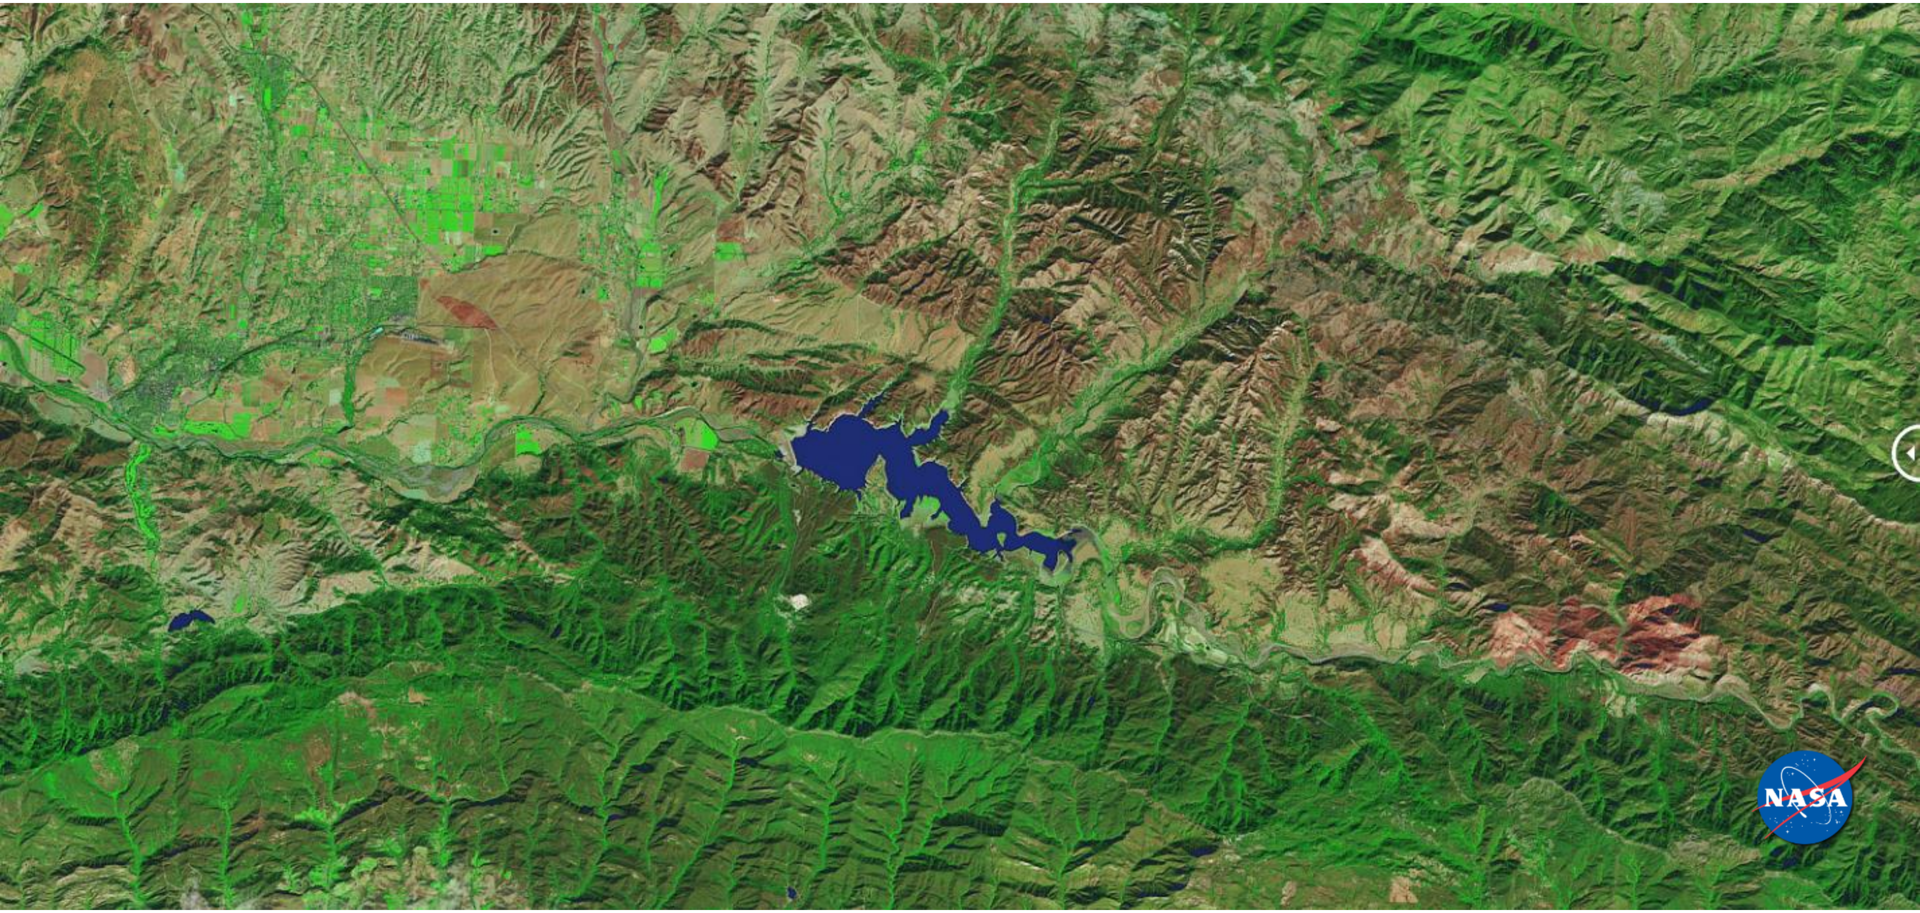


Persistent drought shrinks Lake Cachuma, Southern California

**October 27, 2013** - October 19, 2016

# Motivation

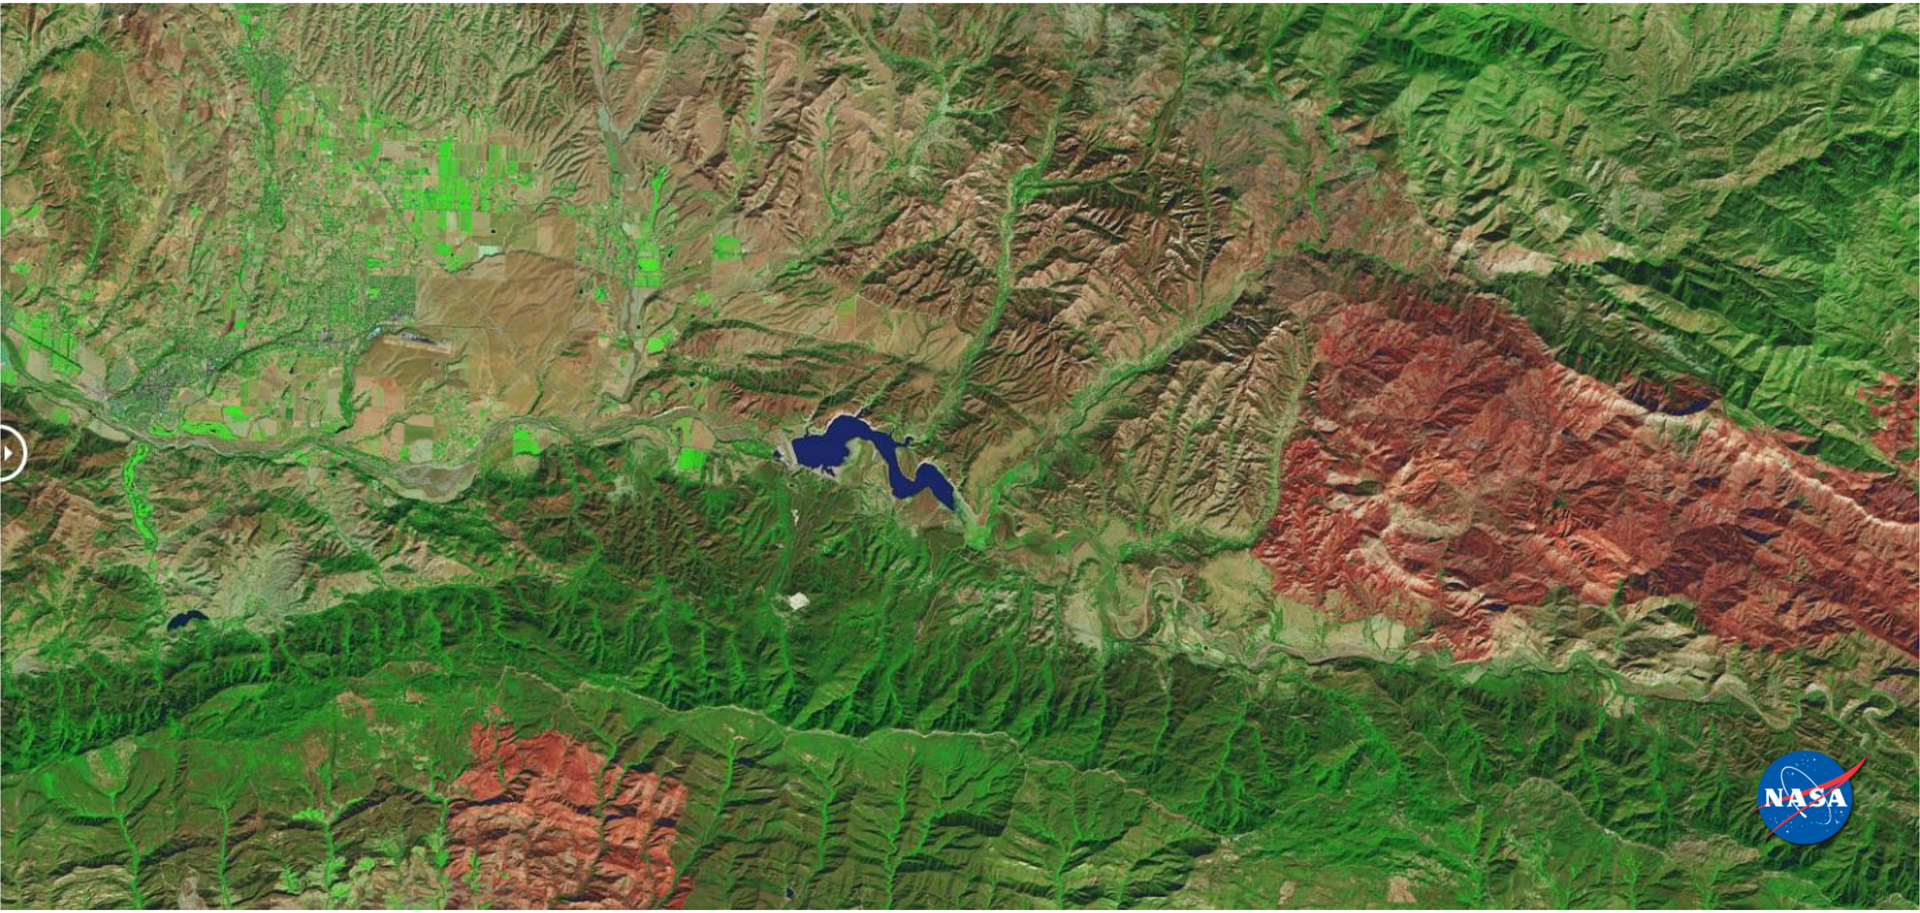


Persistent drought shrinks Lake Cachuma, Southern California
October 27, 2013 - **October 19, 2016**

# Motivation

# Motivation

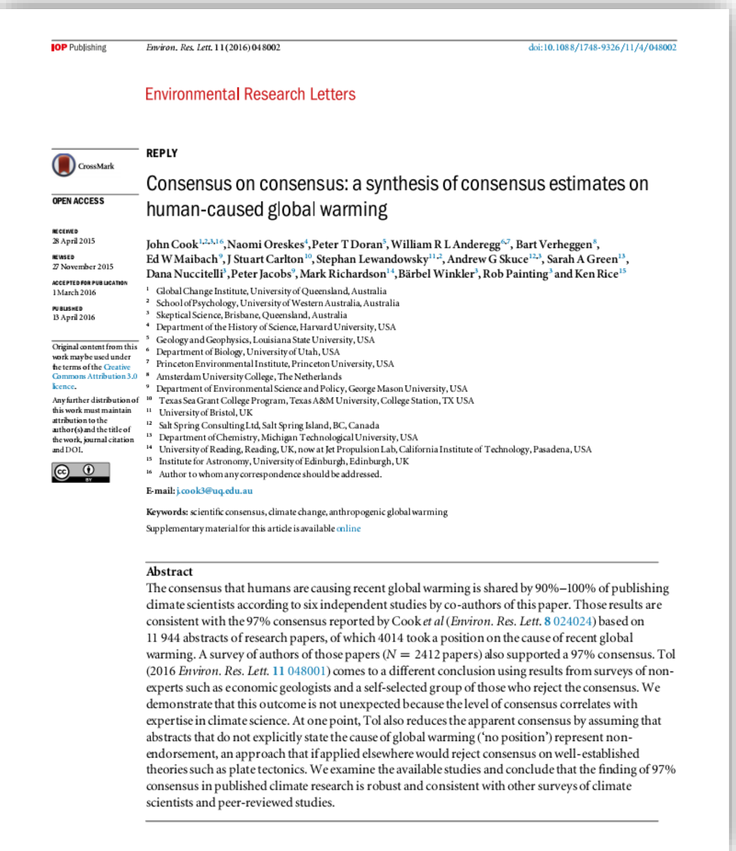
IOP Publishing Environ. Res. Lett. 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

CrossMark

OPEN ACCESS

RECEIVED
28 April 2015

REVISED
27 November 2015

ACCEPTED FOR PUBLICATION
1 March 2016

PUBLISHED
13 April 2016



REPLY

## Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], Naomi Oreskes[4], Peter T Doran[5], William R L Anderegg[6,7], Bart Verheggen[8], Ed W Maibach[9], J Stuart Carlton[10], Stephan Lewandowsky[11,2], Andrew G Skuce[12,3], Sarah A Green[13], Dana Nuccitelli[3], Peter Jacobs[9], Mark Richardson[14], Bärbel Winkler[3], Rob Painting[3] and Ken Rice[15]

1 Global Change Institute, University of Queensland, Australia
2 School of Psychology, University of Western Australia, Australia
3 Skeptical Science, Brisbane, Queensland, Australia
4 Department of the History of Science, Harvard University, USA
5 Geology and Geophysics, Louisiana State University, USA
6 Department of Biology, University of Utah, USA
7 Princeton Environmental Institute, Princeton University, USA
8 Amsterdam University College, The Netherlands
9 Department of Environmental Science and Policy, George Mason University, USA
10 Texas Sea Grant College Program, Texas A&M University, College Station, TX USA
11 University of Bristol, UK
12 Salt Spring Consulting Ltd, Salt Spring Island, BC, Canada
13 Department of Chemistry, Michigan Technological University, USA
14 University of Reading, Reading, UK, now at Jet Propulsion Lab, California Institute of Technology, Pasadena, USA
15 Institute for Astronomy, University of Edinburgh, Edinburgh, UK
16 Author to whom any correspondence should be addressed.

E-mail: j.cook3@uq.edu.au

**Keywords:** scientific consensus, climate change, anthropogenic global warming

Supplementary material for this article is available online

**Abstract**

The consensus that humans are causing recent global warming is shared by 90%–100% of publishing climate scientists according to six independent studies by co-authors of this paper. Those results are consistent with the 97% consensus reported by Cook *et al* (*Environ. Res. Lett.* 8 024024) based on 11 944 abstracts of research papers, of which 4014 took a position on the cause of recent global warming. A survey of authors of those papers ($N = 2412$ papers) also supported a 97% consensus. Tol (2016 *Environ. Res. Lett.* 11 048001) comes to a different conclusion using results from surveys of non-experts such as economic geologists and a self-selected group of those who reject the consensus. We demonstrate that this outcome is not unexpected because the level of consensus correlates with expertise in climate science. At one point, Tol also reduces the apparent consensus by assuming that abstracts that do not explicitly state the cause of global warming ('no position') represent non-endorsement, an approach that if applied elsewhere would reject consensus on well-established theories such as plate tectonics. We examine the available studies and conclude that the finding of 97% consensus in published climate research is robust and consistent with other surveys of climate scientists and peer-reviewed studies.

- Climate is changing.
- The potential impact of that change is huge.
- This would lead to catastrophic consequences for humankind.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Motivation

IOP Publishing Environ. Res. Lett. 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

REPLY

Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], Naomi Oreskes[4], Peter T Doran[5], William R L Anderegg[6,7], Bart Verheggen[8], Ed W Maibach[9], J Stuart Carlton[10], Stephan Lewandowsky[11,2], Andrew G Skuce[12,3], Sarah A Green[13], Dana Nuccitelli[3], Peter Jacobs[9], Mark Richardson[14], Bärbel Winkler[3], Rob Painting[3] and Ken Rice[15]

OPEN ACCESS

RECEIVED 28 April 2015

REVISED 27 November 2015

ACCEPTED FOR PUBLICATION 1 March 2016

PUBLISHED 13 April 2016

## Joint science academies' statement: Global response to climate change

### Climate change is real

There will always be uncertainty in understanding a system as complex as the world's climate. However there is now strong evidence that significant global warming is occurring[1]. The evidence comes from direct measurements of rising surface air temperatures and subsurface ocean temperatures and from phenomena such as increases in average global sea levels, retreating glaciers, and changes to many physical and biological systems. It is likely that most of the warming in recent decades can be attributed to human activities (IPCC 2001)[2]. This warming has already led to changes in the Earth's climate.

The existence of greenhouse gases in the atmosphere is vital to life on Earth – in their absence average temperatures would be about 30 centigrade degrees lower than they are today. But human activities are now causing atmospheric concentrations of greenhouse gases – including carbon dioxide, methane, tropospheric ozone, and nitrous oxide – to rise well above pre-industrial levels. Carbon dioxide levels have increased from 280 ppm in 1750 to over 375 ppm today – higher than any previous levels that can be reliably measured (i.e. in the last 420,000 years). Increasing greenhouse gases are causing temperatures to rise; the Earth's surface warmed by approximately 0.6 centigrade degrees over the twentieth century. The Intergovernmental Panel on Climate Change (IPCC) projected that the average global surface temperatures will continue to increase to between 1.4 centigrade degrees and 5.8 centigrade degrees above 1990 levels, by 2100.

### Reduce the causes of climate change

The scientific understanding of climate change is now sufficiently clear to justify nations taking prompt action. It is vital that all nations identify cost-effective steps that they can take now, to contribute to substantial and long-term reduction in net global greenhouse gas emissions.

Action taken now to reduce significantly the build-up of greenhouse gases in the atmosphere will lessen the magnitude and rate of climate change. As the United Nations Framework Convention on Climate Change (UNFCCC) recognises, a lack of full scientific certainty about some aspects of climate change is not a reason for delaying an immediate response that will, at a reasonable cost, prevent dangerous anthropogenic interference with the climate system.

potentially cost-effective technological options that could contribute to stabilising greenhouse gas concentrations. These are at various stages of research and development. However barriers to their broad deployment still need to be overcome.

Carbon dioxide can remain in the atmosphere for many decades. Even with possible lowered emission rates we will be experiencing the impacts of climate change throughout the 21st century and beyond. Failure to implement significant reductions in net greenhouse gas emissions now, will make the job much harder in the future.

### Prepare for the consequences of climate change

Major parts of the climate system respond slowly to changes in greenhouse gas concentrations. Even if greenhouse gas emissions were stabilised instantly at today's levels, the climate would still continue to change as it adapts to the increased emission of recent decades. Further changes in climate are therefore unavoidable. Nations must prepare for them.

The projected changes in climate will have both beneficial and adverse effects at the regional level, for example on water resources, agriculture, natural ecosystems and human health. The larger and faster the changes in climate, the more likely it is that adverse effects will dominate. Increasing temperatures are likely to increase the frequency and severity of weather events such as heat waves and heavy rainfall. Increasing temperatures could lead to large-scale effects such as melting of large ice sheets (with major impacts on low-lying regions throughout the world). The IPCC estimates that the combined effects of ice melting and sea water expansion from ocean warming are projected to cause the global mean sea-level to rise by between 0.1 and 0.9 metres between 1990 and 2100. In Bangladesh alone, a 0.5 metre sea-level rise would place about 6 million people at risk from flooding.

Developing nations that lack the infrastructure or resources to respond to the impacts of climate change will be particularly affected. It is clear that many of the world's poorest people are likely to suffer the most from climate change. Long-term global efforts to create a more healthy, prosperous and sustainable world may be severely hindered by changes in the climate.

- Climate is changing.
- The potential impact of that change is huge.
- This would lead to catastrophic consequences for humankind.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos UAB Universitat Autònoma de Barcelona

# Motivation

IOP Publishing *Environ. Res. Lett.* 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

CrossMark

OPEN ACCESS

RECEIVED
28 April 2015

REVISED
27 November 2015

ACCEPTED FOR PUBLICATION
1 March 2016

PUBLISHED
13 April 2016

REPLY

Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], Naomi Oreskes[4], Peter T Doran[5], William R L Anderegg[6,7], Bart Verheggen[8], Ed W Maibach[9], J Stuart Carlton[10], Stephan Lewandowsky[11,2], Andrew G Skuce[12,3], Sarah A Green[13], Dana Nuccitelli[3], Peter Jacobs[9], Mark Richardson[14], Bärbel Winkler[3], Rob Painting[3] and Ken Rice[15]

Joint science academies' statement: Global response to climate change

**Climate change is real**

There will always be uncertainty in understanding a system as complex as the world's climate. However there is now strong evidence that significant global warming is occurring[1]. The evidence comes from direct measurements of rising surface air temperatures and subsurface ocean temperatures and from phenomena such as increases in average global sea levels, retreating glaciers, and changes to many physical and biological systems. It is likely that most of the warming in recent decades can be attributed to human activities (IPCC 2001)[2]. This warming has already led to changes in the Earth's climate.

The existence of greenhouse gases in the atmosphere is vital to life on Earth – in their absence average temperatures would be about 30 centigrade degrees lower than they are today. But human activities are now causing atmospheric concentrations of greenhouse gases – including carbon dioxide, methane, tropospheric ozone, and nitrous oxide – to rise well above pre-industrial levels. Carbon dioxide levels have increased from 280 ppm in 1750 to over 375 ppm today – higher than any previous levels that can be reliably measured (i.e. in the last 420,000 years). Increasing greenhouse gases are causing temperatures to rise; the Earth's surface warmed by approximately 0.6 centigrade degrees over the twentieth century. The Intergovernmental Panel on Climate Change (IPCC) projected that the average global surface temperatures will continue to increase to between 1.4 centigrade degrees and 5.8 centigrade degrees above 1990 levels, by 2100.

**Reduce the causes of climate change**

The scientific understanding of climate change is now sufficiently clear to justify nations taking prompt action. It is vital that all nations identify cost-effective steps that they can take now, to contribute to substantial and long-term reduction in net global greenhouse gas emissions.

Action taken now to reduce significantly the build-up of greenhouse gases in the atmosphere will lessen the magnitude and rate of climate change. As the United Nations Framework Convention on Climate Change (UNFCCC) recognises, a lack of full scientific certainty about some aspects of climate change is not a reason for delaying an immediate response that will, at a reasonable cost, prevent dangerous anthropogenic interference with the climate system.

potentially cost-effective technological options that could contribute to stabilising greenhouse gas concentrations. These are at various stages of research and development. However barriers to their broad deployment still need to be overcome.

Carbon dioxide can remain in the atmosphere for many decades. Even with possible lowered emission rates we will be experiencing the impacts of climate change throughout the 21st century and beyond. Failure to implement significant reductions in net greenhouse gas emissions now, will make the job much harder in the future.

**Prepare for the consequences of climate change**

Major parts of the climate system respond slowly to changes in greenhouse gas concentrations. Even if greenhouse gas emissions were stabilised instantly at today's levels, the climate would still continue to change as it adapts to the increased emission of recent decades. Further changes in climate are therefore unavoidable. Nations must prepare for them.

The projected changes in climate will have both beneficial and adverse effects at the regional level, for example on water resources, agriculture, natural ecosystems and human health. The larger and faster the changes in climate, the more likely it is that adverse effects will dominate. Increasing temperatures are likely to increase the frequency and severity of weather events such as heat waves and heavy rainfall. Increasing temperatures could lead to large-scale effects such as melting of large ice sheets (with major impacts on low-lying regions throughout the world). The IPCC estimates that the combined effects of ice melting and sea water expansion from ocean warming are projected to cause the global mean sea-level to rise by between 0.1 and 0.9 metres between 1990 and 2100. In Bangladesh alone, a 0.5 metre sea-level rise would place about 6 million people at risk from flooding.

Developing nations that lack the infrastructure or resources to respond to the impacts of climate change will be particularly affected. It is clear that many of the world's poorest people are likely to suffer the most from climate change. Long-term global efforts to create a more healthy, prosperous and sustainable world may be severely hindered by changes in the climate.

- Climate is changing.
- The potential impact of that change is huge.
- This would lead to catastrophic consequences for humankind.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

# Motivation


IOP Publishing Environ. Res. Lett. 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

OPEN ACCESS

RECEIVED 28 April 2015

REVISED 27 November 2015

PUBLISHED 13 April 2016

REPLY

Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], Naomi Oreskes[4], Peter T Doran[5], William R L Anderegg[6,7], Bart Verheggen[8], Ed W Maibach[9], J Stuart Carlton[10], Stephan Lewandowsky[11,2], Andrew G Skuce[12,3], Sarah A Green[13], Dana Nuccitelli[3], Peter Jacobs[9], Mark Richardson[14], Bärbel Winkler[3], Rob Painting[3] and Ken Rice[15]

Joint science academies' statement: Global response to climate change

**Climate change is real**

There will always be uncertainty in understanding a system as complex as the world's climate. However there is now strong evidence that significant global warming is occurring[1]. The evidence comes from direct measurements of rising surface air temperatures and subsurface ocean temperatures and from phenomena such as increases in average global sea levels, retreating glaciers, and changes to many physical and biological systems. It is likely that most of the warming in recent decades can be attributed to human activities (IPCC 2001)[2]. This warming has already led to changes in the Earth's climate.

The existence of greenhouse gases in the atmosphere is vital to life on Earth – in their absence average temperatures would be about 30 centigrade degrees than they are today. But human activities are now cau atmospheric concentrations of greenhouse gases – including carbon dioxide, methane, tropospheric ozo and nitrous oxide – to rise well above pre-industrial le Carbon dioxide levels have increased from 280 ppm i 1750 to over 375 ppm today – higher than any previo levels that can be reliably measured (i.e. in the last 42 years). Increasing greenhouse gases are causing temperatures to rise; the Earth's surface warmed by approximately 0.6 centigrade degrees over the twenti century. The Intergovernmental Panel on Climate Cha (IPCC) projected that the average global surface temperatures will continue to increase to between 1.4 centigrade degrees and 5.8 centigrade degrees above levels, by 2100.

**Reduce the causes of climate change**

The scientific understanding of climate change is now sufficiently clear to justify nations taking prompt actio is vital that all nations identify cost-effective steps tha can take now, to contribute to substantial and long-te reduction in net global greenhouse gas emissions.

Action taken now to reduce significantly the build-up greenhouse gases in the atmosphere will lessen the magnitude and rate of climate change. As the United Nations Framework Convention on Climate Change (UNFCCC) recognises, a lack of full scientific certainty about some aspects of climate change is not a reason delaying an immediate response that will, at a reason cost, prevent dangerous anthropogenic interference w the climate system.

potentially cost-effective technological options that could contribute to stabilising greenhouse gas concentrations. These are at various stages of research and development. However barriers to their broad deployment still need to be overcome.

Carbon dioxide can remain in the atmosphere for many decades. Even with possible lowered emission rates we will be experiencing the impacts of climate change throughout the 21st century and beyond. Failure to implement significant reductions in net greenhouse gas emissions now, will make the job much harder in the future.

**Prepare for the consequences of**

- Climate is changing.
- The potential impact of that change is huge.
- This would lead to catastrophic consequences for humankind.

# Motivation

# Motivation

- Projections
- Impact analysis
- Adaptation to climate change.

# Motivation

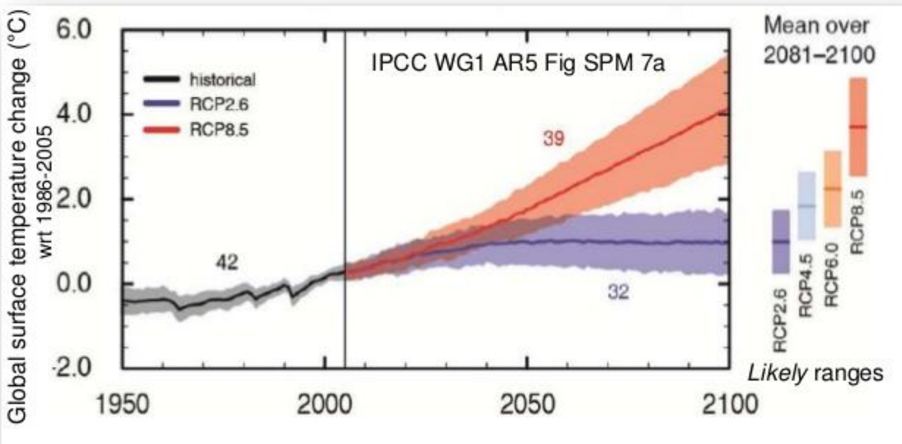


- Projections
- Impact analysis
- Adaptation to climate change.

# Motivation

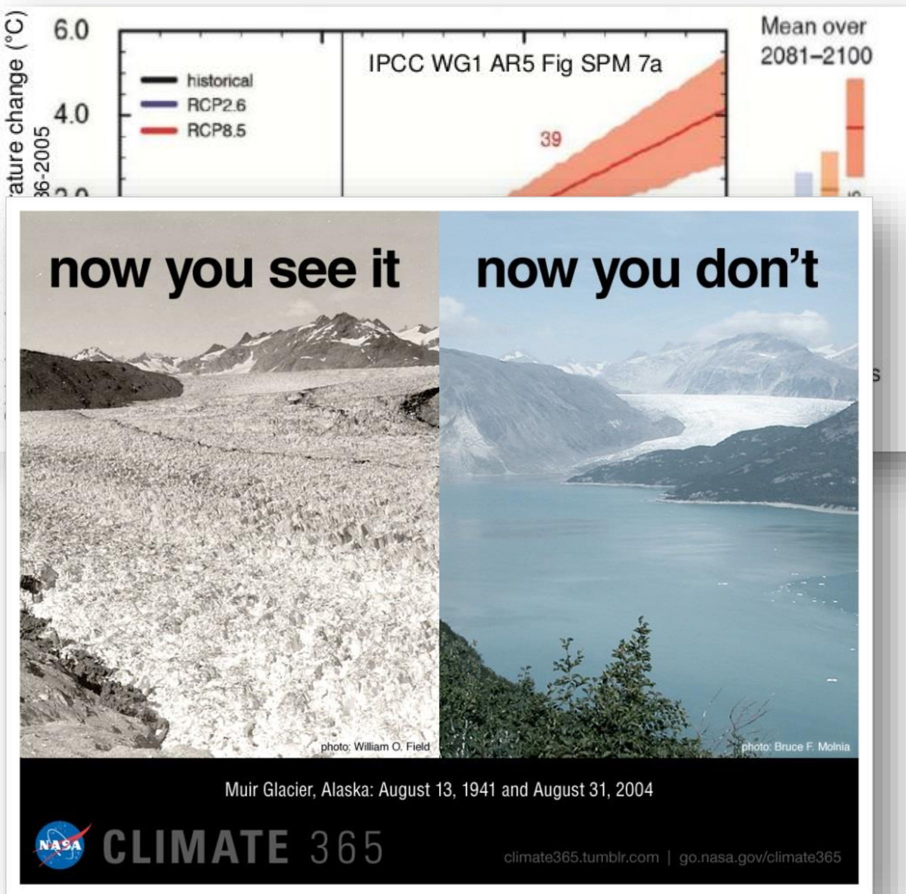


- Projections
- Impact analysis
- Adaptation to climate change.

# Motivation


- Projections
- Impact analysis
- Adaptation to climate change.

# Motivation

- Projections
- Impact analysis
- Adaptation to climate change.

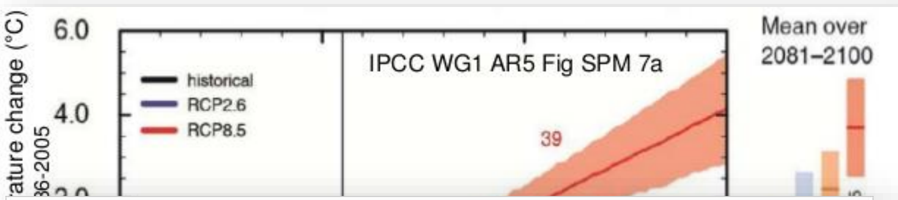



Currently, **only computational models have the potential** to provide geographically and physically consistent estimates.

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

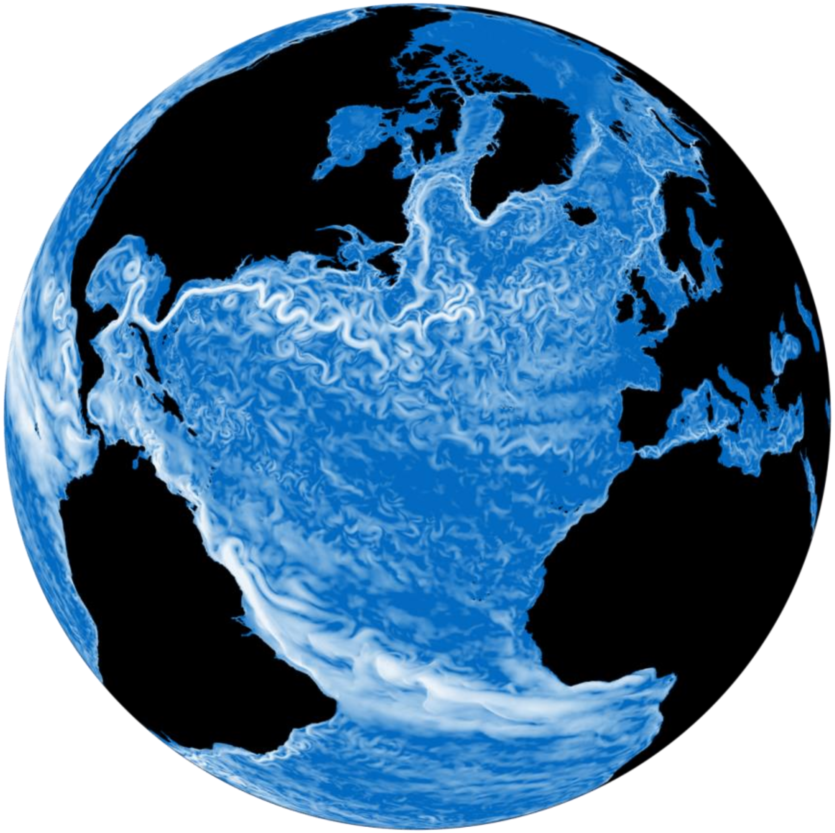

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

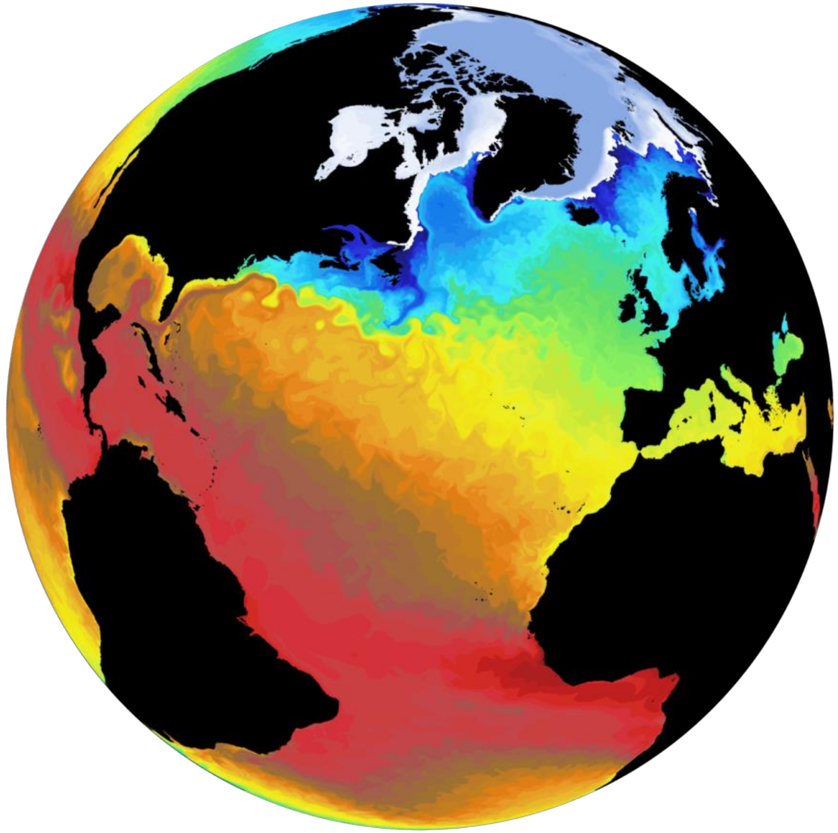

Current Velocity

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

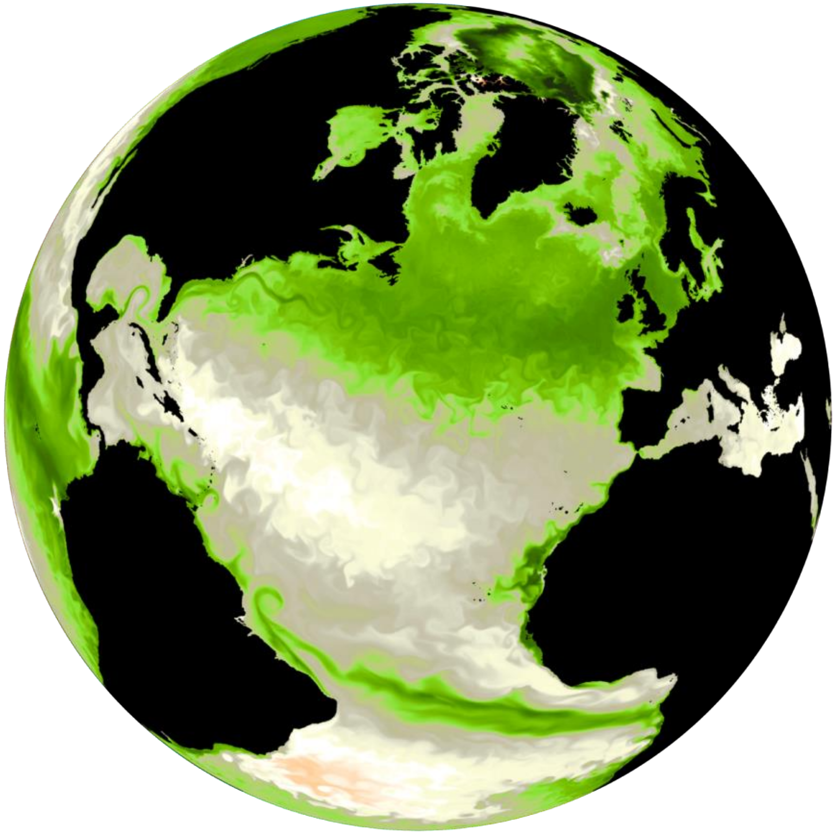

Sea Surface Temperature

Current Velocity

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# NEMO

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

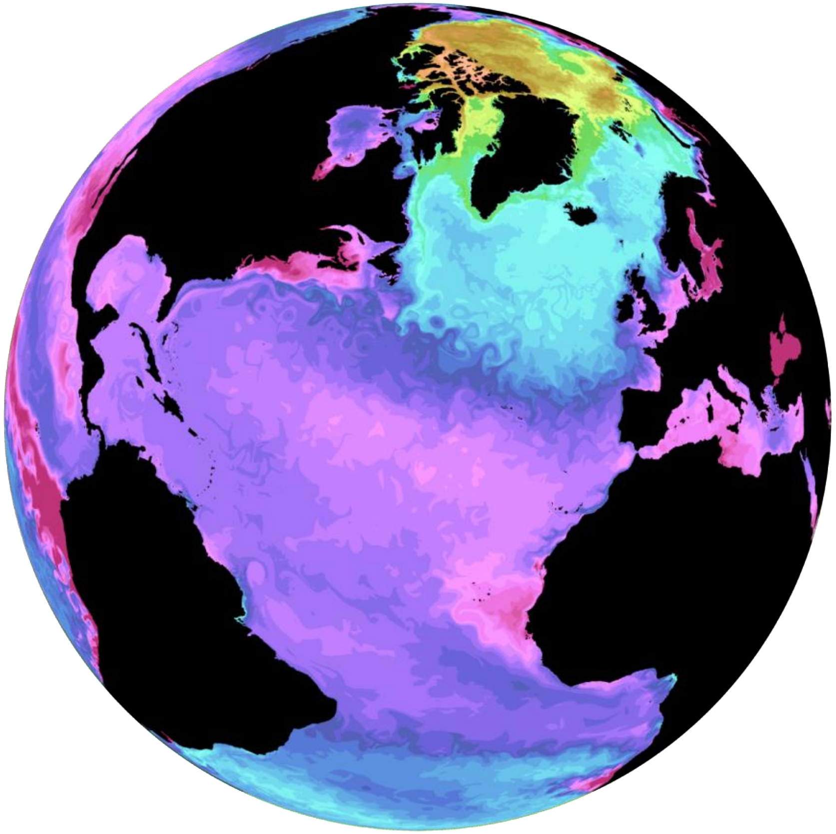

Sea Surface Temperature

Current Velocity

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

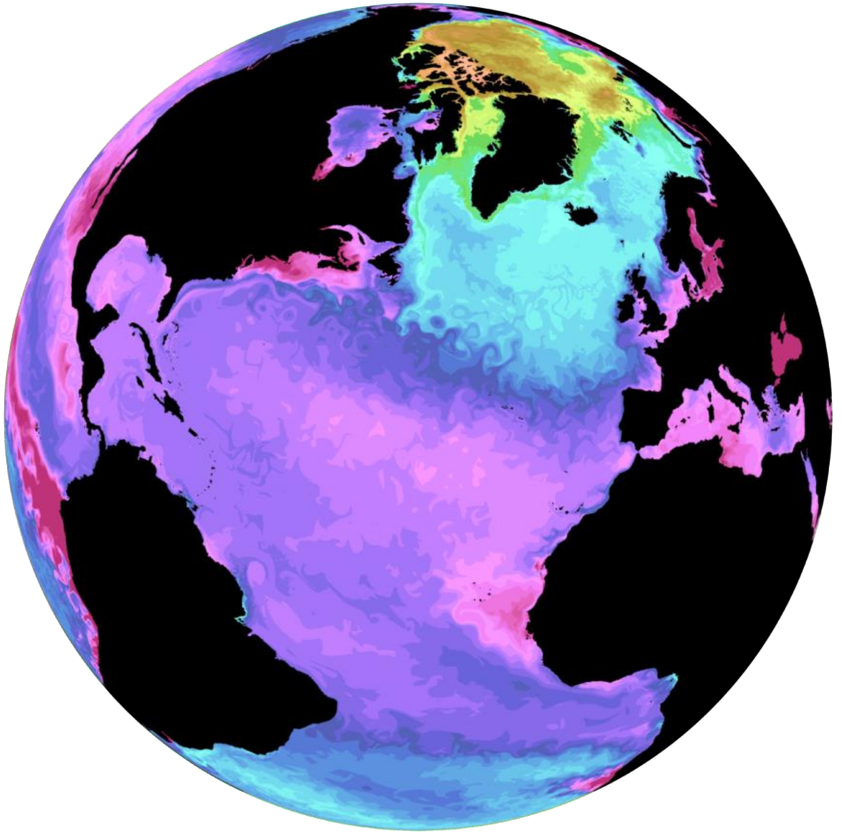

Sea Surface Temperature

Chlorophyll

Velocity

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

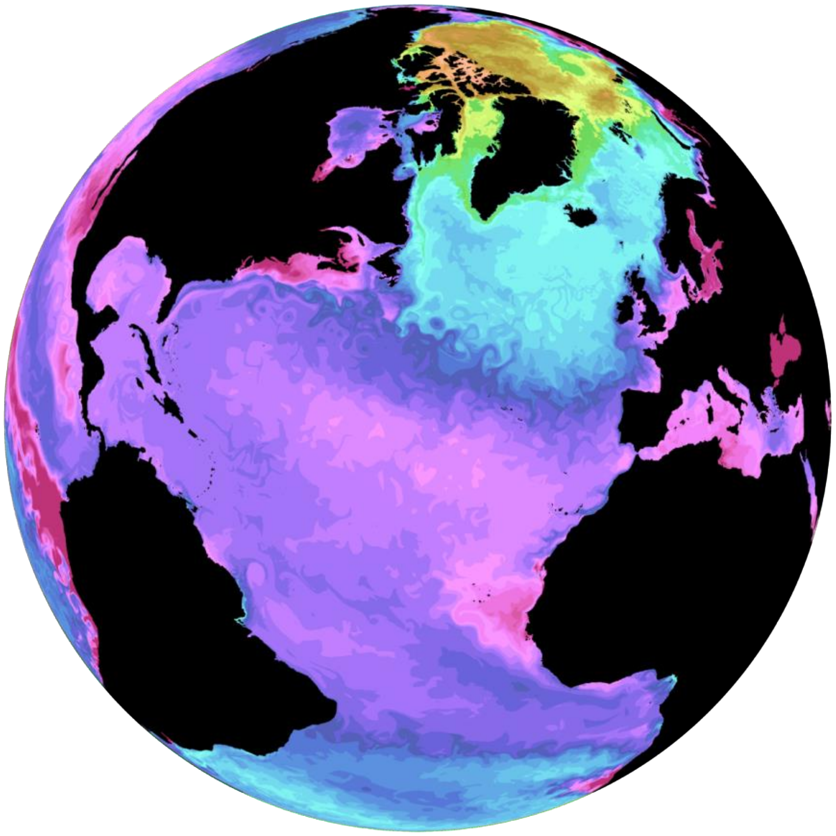

Sea Surface Temperature
Currents Velocity
Chlorophyll

NEMO

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

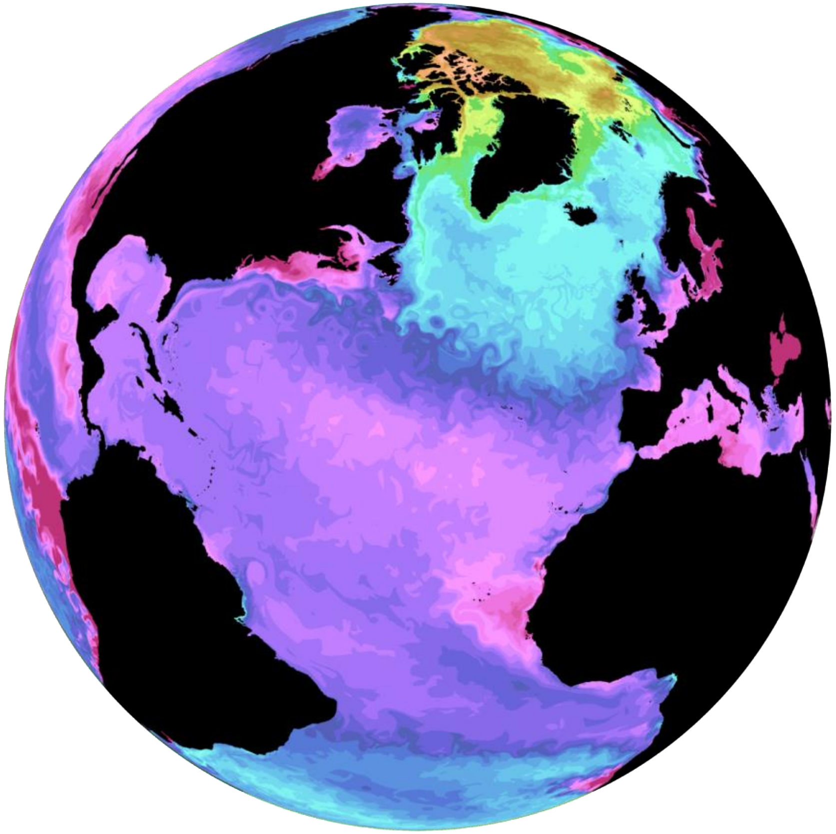

Sea Surface Temperature

Chlorophyll

Velocity

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

**Nucleus for European Modeling of the Ocean (NEMO)** is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

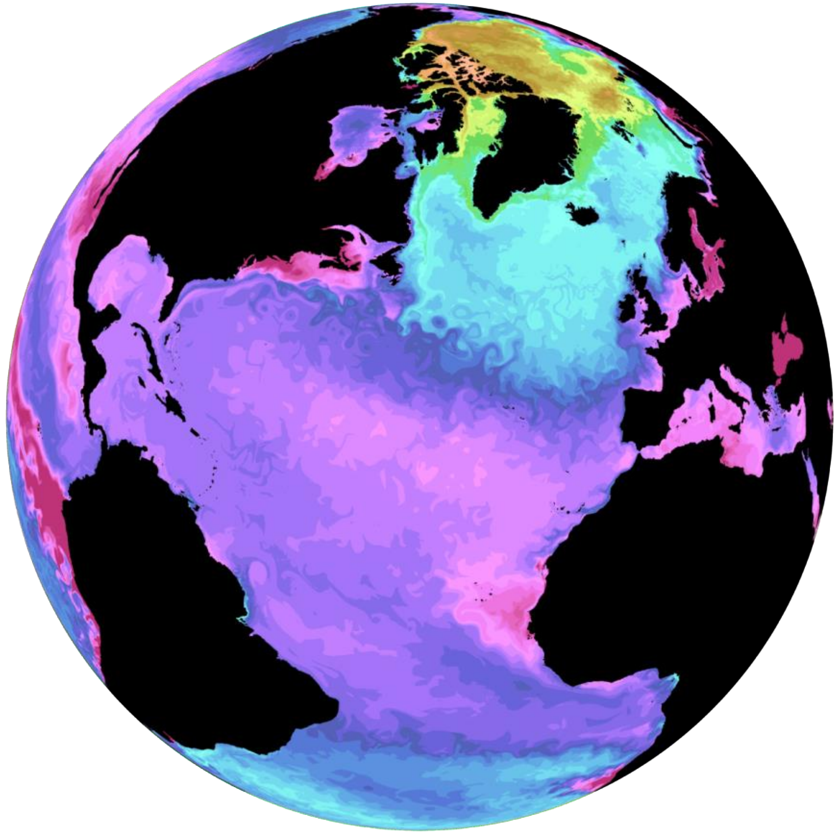

Sea Surface Temperature

Currents Velocity

Chlorophyll

# A climate modeling Timeline
## Inclusion of new components

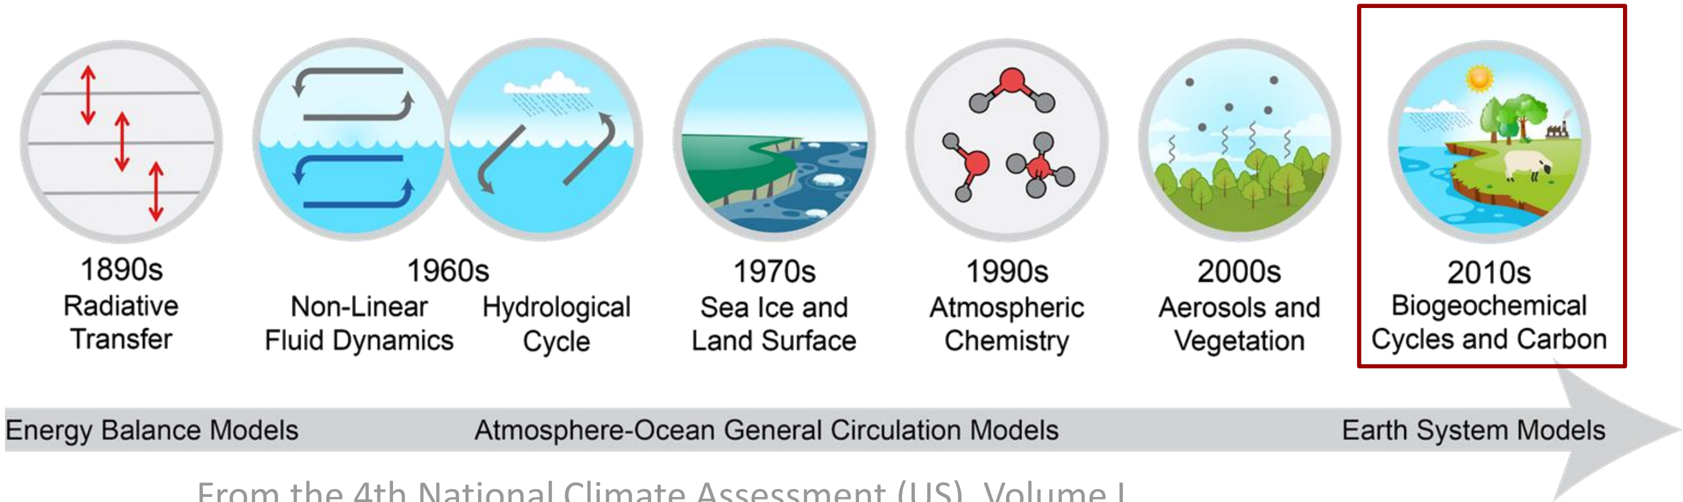


From the 4th National Climate Assessment (US), Volume I

- Allowed the representation of new climate and biogeochemical processes
- Improved the ESMs ability to represent the real world
- Provides a new framework to investigate the interactions between the different components

# A climate modeling Timeline

## Increase in spatial resolution: Atmosphere

Typical climate model in 90s

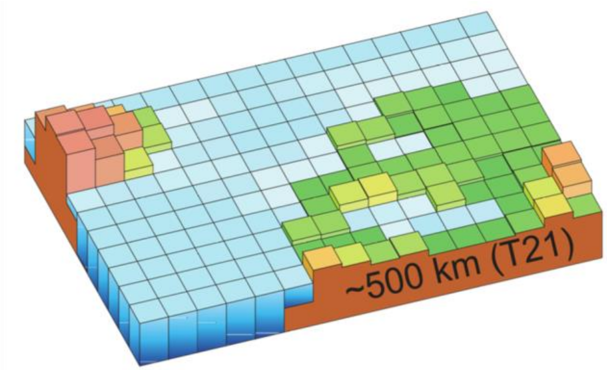


Today

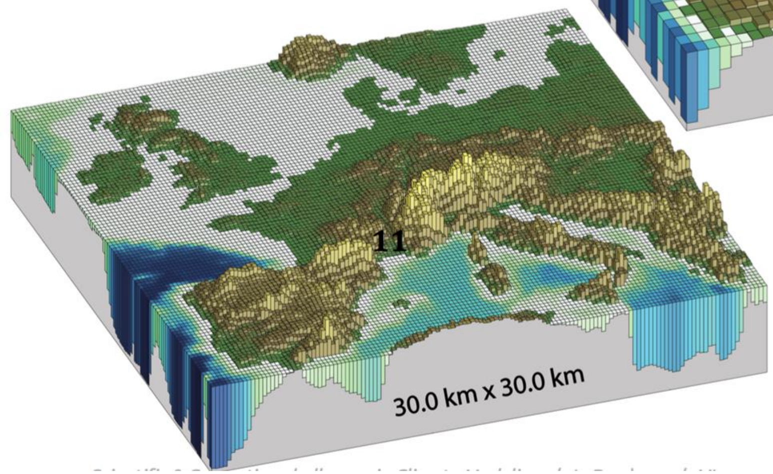


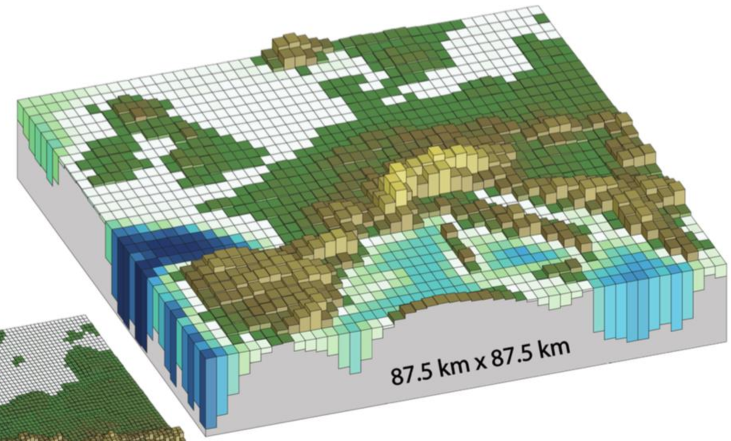


Default resolution

“High” resolution

Achieving higher resolutions is essential to better represent orography, and its effect on climate (i.e. in precipitation)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# A climate modeling Timeline
## Increase in spatial resolution: Ocean

Standard Resolution (1°)

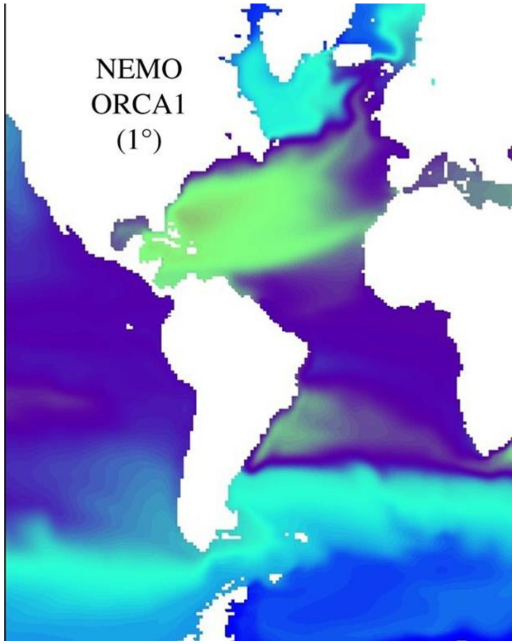


High Resolution (0.25°) Eddy-permitting

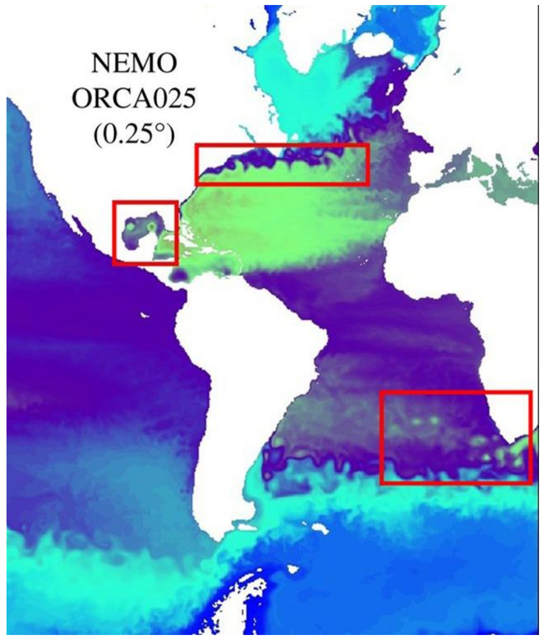


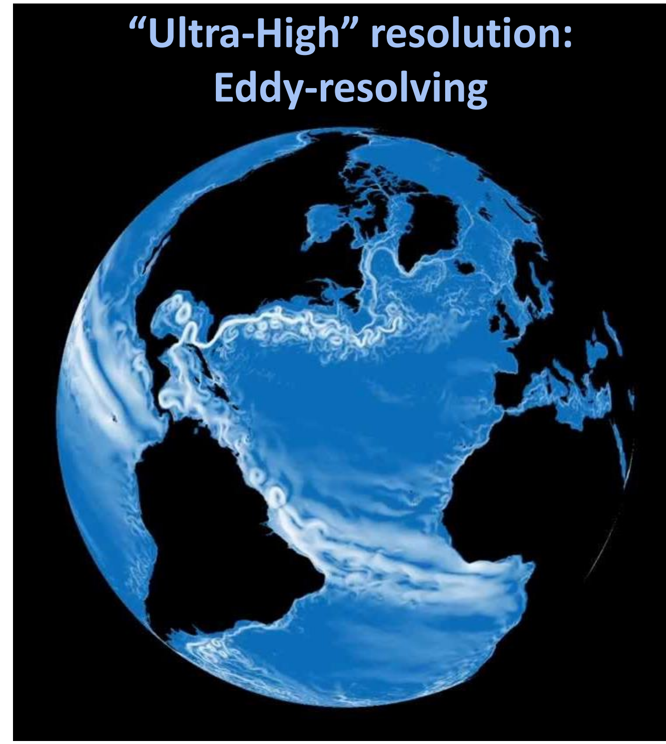


The improvements in ocean resolution translate in a better representation of eddies and ocean currents, which are key to describe realistically decadal variability in the ocean

# Challenges

# Challenges

TOP500 Ranking of most powerful supercomputers

# Challenges

TOP500 Ranking of most powerful supercomputers

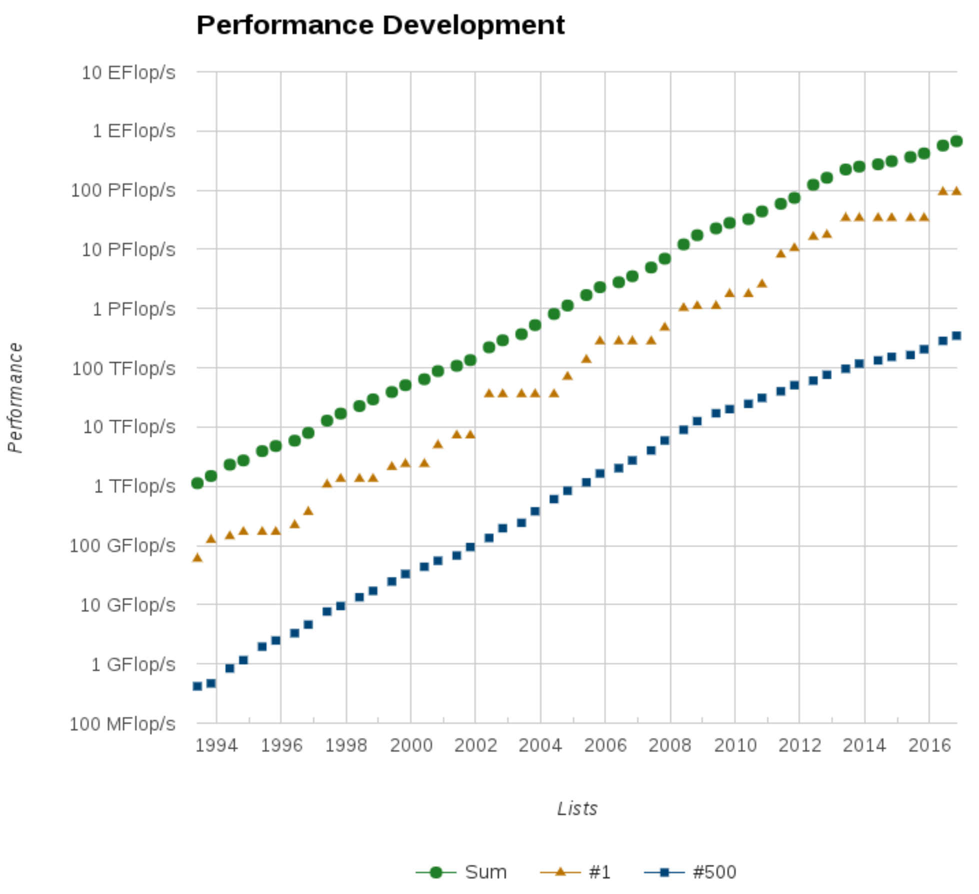

# Challenges

35 YEARS OF MICROPROCESSOR TREND DATA

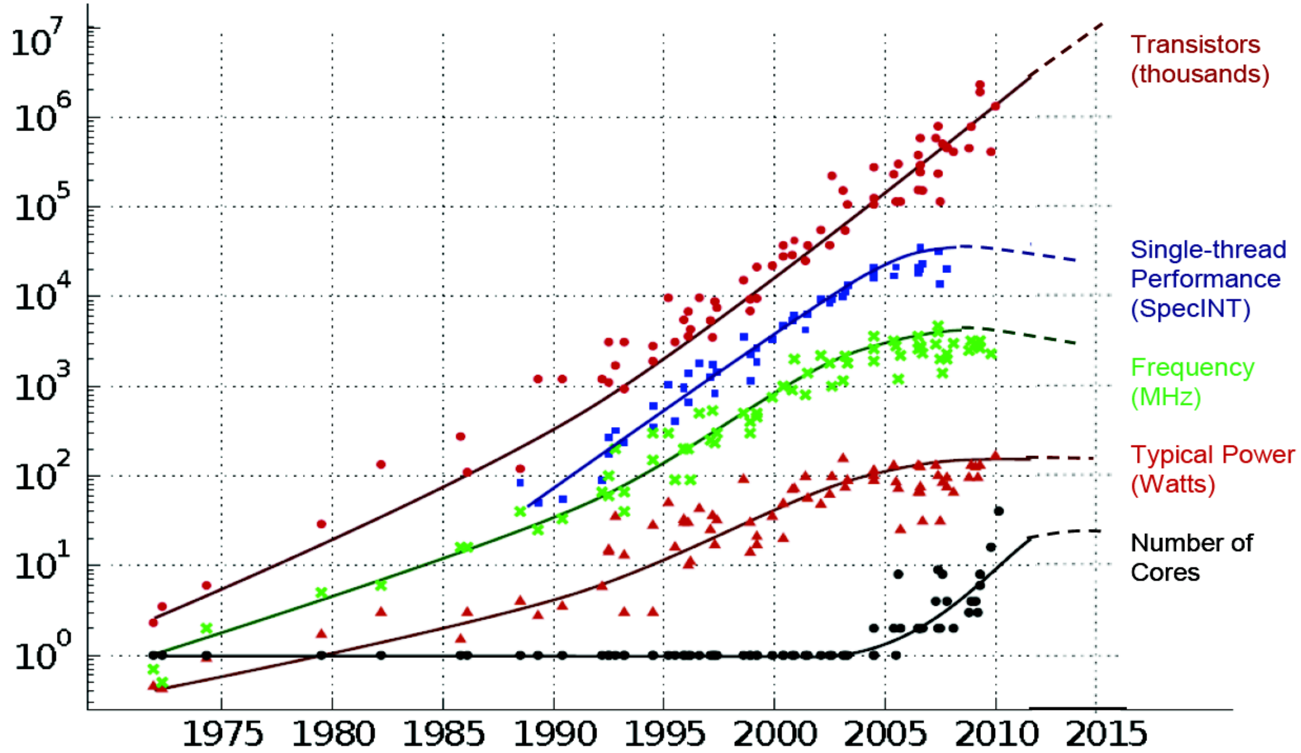


Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten
Dotted line extrapolations by C. Moore

# Challenges

# Challenges

- To be able to use the computing power of modern supercomputers, applications must exploit extreme parallelism. To do so, applications must use:
  - Synchonization-reducing algorithms.
  - Communication reducing algorithms.
  - Mixed Precision methods.
  - Autotuning.
  - Fault resilient algorithms.

# Challenges

- Start from scratch to use better the resources?

# Challenges

- Start from scratch to use better the resources?

  Development cost of a climate model is estimated to be between 500-1000 person year.

# Challenges

- Start from scratch to use better the resources?

  Development cost of a climate model is estimated to be between 500-1000 person year.

  Most of models are community models and communities are reluctant to deep changes

# Challenges

- Start from scratch to use better the resources?

  Development cost of a climate model is estimated to be between 500-1000 person year.

  Most of models are community models and communities are reluctant to deep changes

- Adapt the existing state-of-the-art models seems wiser choice!

# Objectives

- Perform a computational performance analysis of the **NEMO** model
- Suggest and apply optimizations to increase model's performance on modern supercomputers

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Work Done

# Stress Test

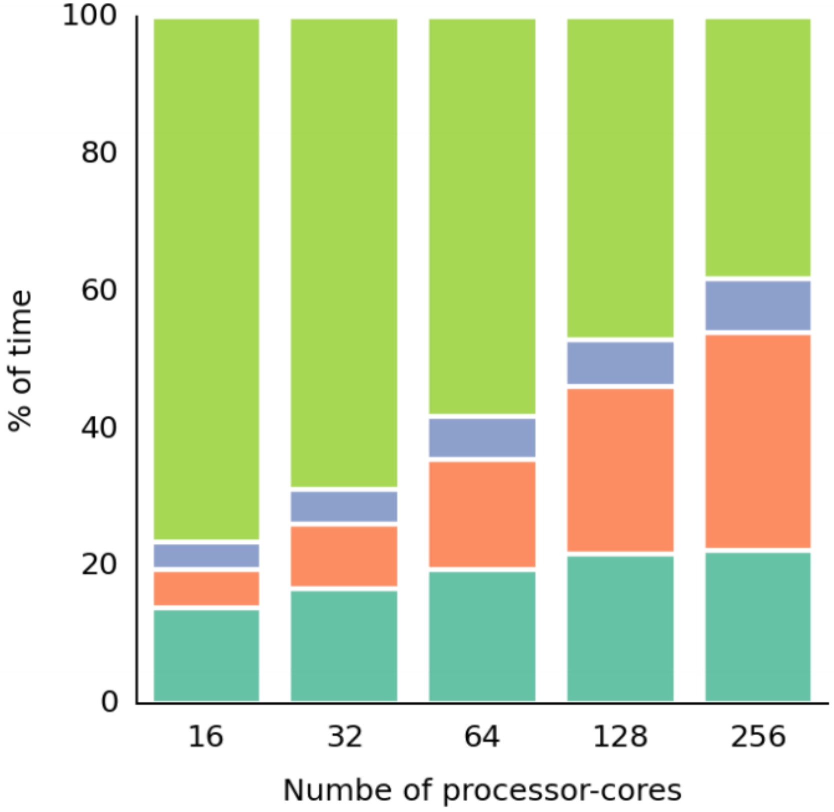


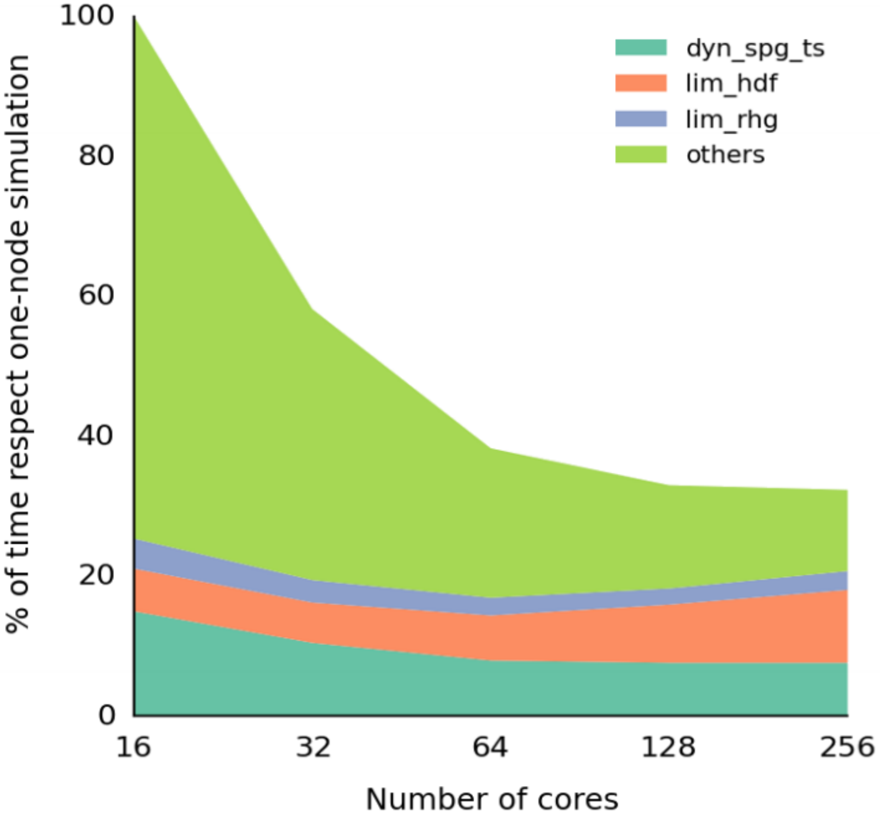


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Impact of Communication

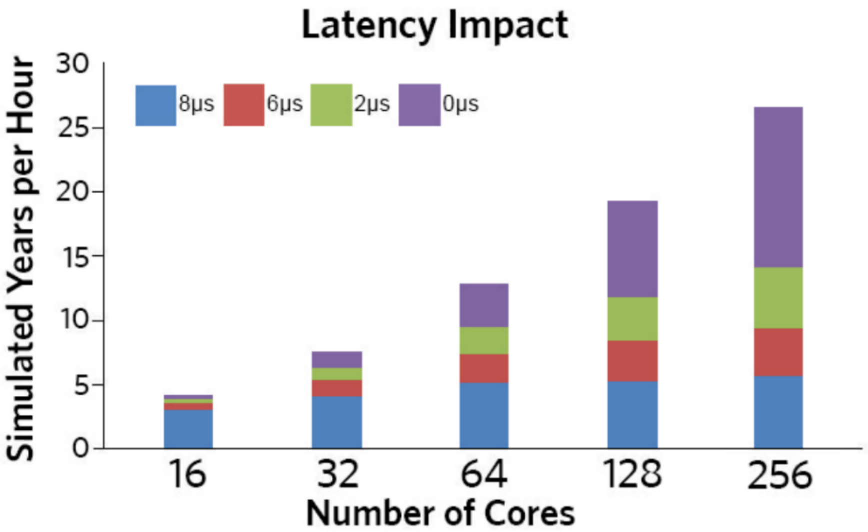

# Optimizations

- Reduction of communication:
  - Message agregation
  - Removing unnecessary collective communication

# Optimizations

## MESSAGE PACKING


Taking in account that NEMO is much more sensitive to latency than to bandwidth, the aggregation of messages is the best way to reduce the time invested in communications. Therefore, consecutive messages can be packed wherever the computational dependencies allow to do so. This approach can be applied both to point to point and collective communication.

# Optimizations

## CONVERGENCE CHECK REDUCTION

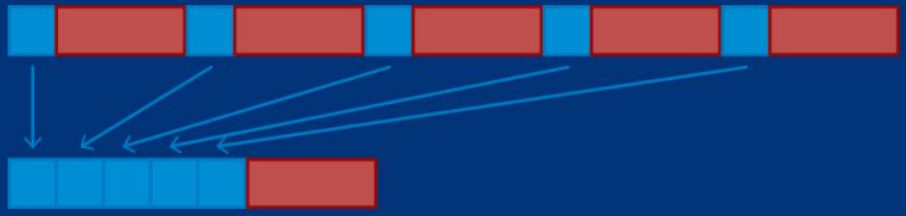

Some routines use collective communications to perform a convergence check in iterative solvers. The cost of this verifications is really high, reaching a 66% of the time in some cases. Wherever the model allowed it, we reduced the frequency of this verifications in order to increase the parallel efficiency.

# Optimizations

## REORDERING

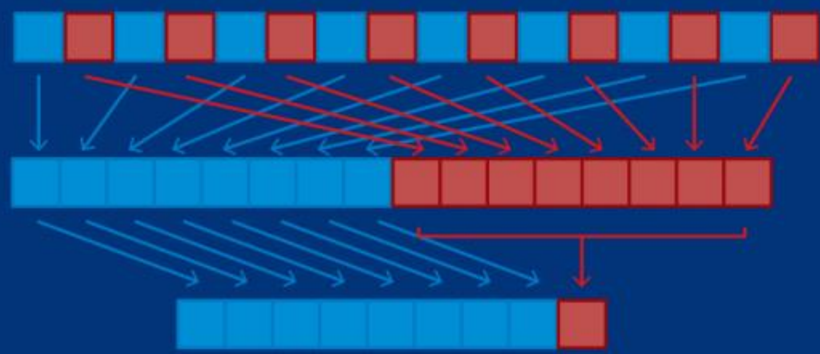

In order to apply the message packing optimization to as many routines as possible, it was necessary to rearrange some computation and communication regions, taking into account the dependencies between them, to reduce the number of messages.

# Optimization Impact

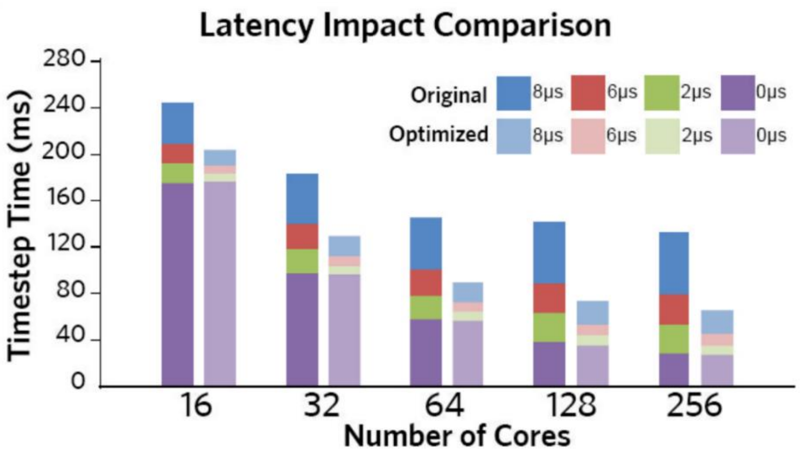

# Optimization Impact on ORCA025 simulations

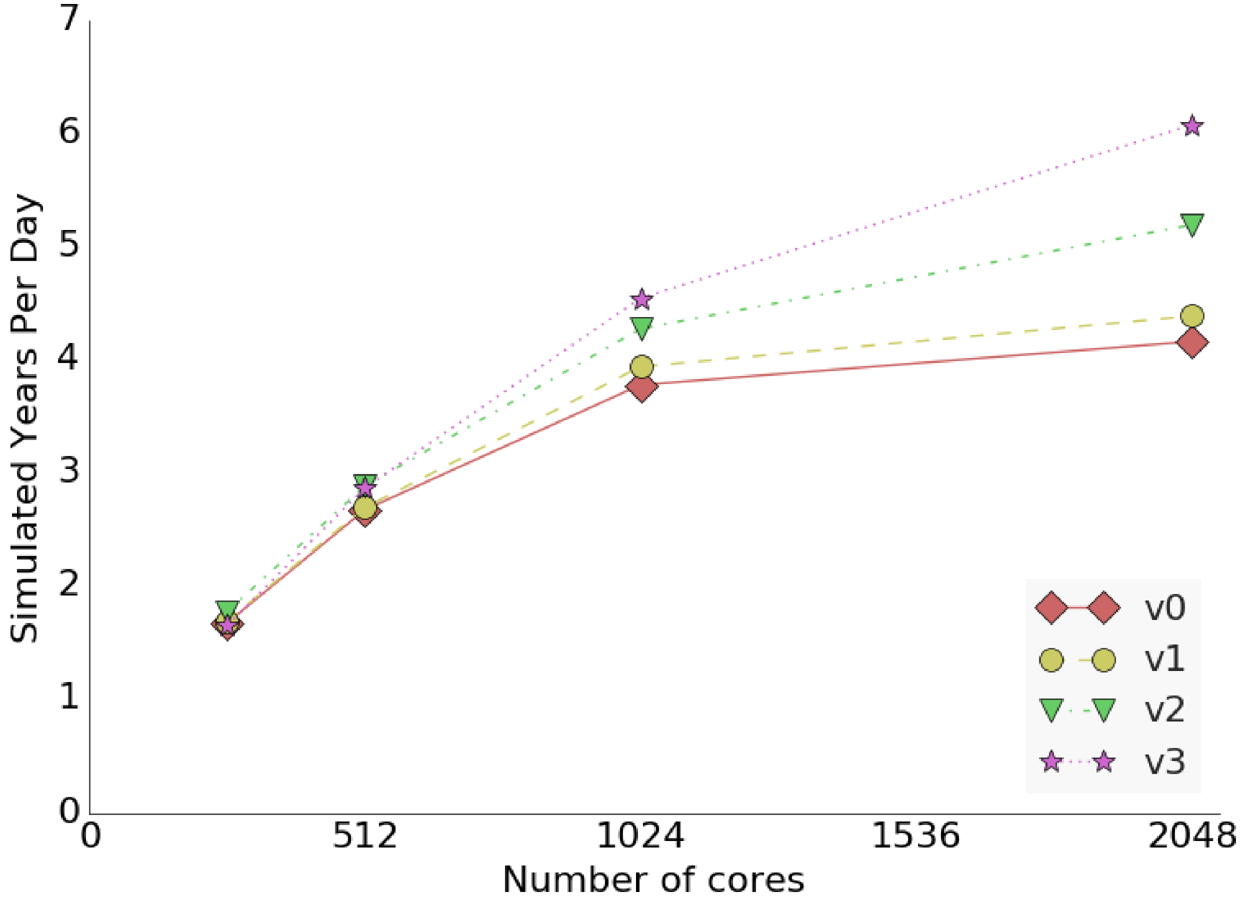


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
UAB Universitat Autònoma de Barcelona

# Impact of our work

- Our optimizations have been included in the latests stable release. Hundreds of users arround the world will take profit of our it and millions of computer hours can be saved.

- Paper under revision with the methodology used to analyze the model, that can be useful for other Earth Science models.

Finding, analyzing and solving MPI communication bottlenecks in Earth System models

Oriol Tintó Prims[1,2], Miguel Castrillo[1], Mario C. Acosta[1], Oriol Mula-Valls[1], Kim Serradell[1], Ana Cortés[2], Francisco Doblas-Reyes[1,3]

[1] *Barcelona Supercomputing Center*
[2] *Universitat Autònoma de Barcelona*
[3] *Catalan Institution for Research and Advanced Studies, ICREA*

**Abstract**

It is a matter of consensus that the ability to efficiently use current and future High Performance Computing systems is crucial for science, however, the reality is that the performance currently achieved by most of the parallel scientific applications is far from desired. Despite inter-process communication has already been a matter of study in many different works, it is a fact that their recommendations are not taken into account in most of computational model development processes, at least in the case of Earth Science. This work presents a methodology that aims to help scientists working with computational models using inter-process communication to deal with the difficulties they face when trying to understand their applications behaviour. Following a series of steps that are presented here, both users and developers will learn how to identify performance issues by characterizing applications scalability, identifying which parts present a bad performance and understand the role that inter-process communication plays. In this work, the Nucleus for European Modelling of the Ocean (NEMO), the state-of-the-art European global ocean circulation model, will be used as an example of success. It is a community code widely used in Europe, to the extent that more than a hundred million core hours are used every year in experiments involving NEMO. In the analysis exercise it is shown how to answer the questions of where, why and what is degrading model's scalability, and how this information can help developers in finding solutions that will mitigate their eventual issues. This document also demonstrates how performance analysis carried out with small size experiments using limited resources can lead to optimizations that will impact bigger experiments running on thousands of

*Preprint submitted to Journal of Computational Science* *March 10, 2017*

# Impact of our work

- Our optimizations have been included in the latests stable release. Hundreds of users arround the world will take profit of our it and millions of computer hours can be saved.

- Paper under revision with the methodology used to analyze the model, that can be useful f... other Earth Science m...

ACCEPTED!

Finding, analyzing and solving MPI communication bottlenecks in Earth System models

Oriol Tintó Prims[1,2], Miguel Castrillo[1], Mario C. Acosta[1], Oriol Mula-Valls[1], Kim Serradell[1], Ana Cortés[2], Francisco Doblas-Reyes[1,3]

[1] Barcelona Supercomputing Center
[2] Universitat Autònoma de Barcelona
[3] Catalan Institution for Research and Advanced Studies, ICREA

**Abstract**

It is a matter of consensus that the ability to efficiently use current and future High Performance Computing systems is crucial for science, however, the reality is that the performance currently achieved by most of the parallel scientific applications is far from desired. Despite inter-process communication has already been a matter of study in many different works, it is a fact that their recommendations are not taken into account in most of computational model development processes, at least in the case of Earth Science. This work presents a methodology that aims to help scientists working with computational models using inter-process communication to deal with the difficulties they face when trying to understand their applications behaviour. Following a series of steps that are presented here, both users and developers will learn how to identify performance issues by characterizing applications scalability, identifying which parts present a bad performance and understand the role that inter-process communication plays. In this work, the Nucleus for European Modelling of the Ocean (NEMO), the state-of-the-art European global ocean circulation model, will be used as an example of success. It is a community code widely used in Europe, to the extent that more than a hundred million core hours are used every year in experiments involving NEMO. In the analysis exercise it is shown how to answer the questions of where, why and what is degrading model's scalability, and how this information can help developers in finding solutions that will mitigate their eventual issues. This document also demonstrates how performance analysis carried out with small size experiments using limited resources can lead to optimizations that will impact bigger experiments running on thousands of

*Preprint submitted to Journal of Computational Science* *March 10, 2017*

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

# Impact of our work

- Our optimizations have been included in the latests stable release. Hundreds of users arround the world will take profit of our it and millions of computer hours can be saved.

- Paper under revision with the methodology used to analyze the model, that can be useful f r other Earth Science m

ACCEPTED!

@Journal of Computational Science

Finding, analyzing and solving MPI communication bottlenecks in Earth System models

Oriol Tintó Prims[1,2], Miguel Castrillo[1], Mario C. Acosta[1], Oriol Mula-Valls[1], Kim Serradell[1], Ana Cortés[2], Francisco Doblas-Reyes[1,3]

[1] Barcelona Supercomputing Center
[2] Universitat Autònoma de Barcelona
[3] Catalan Institution for Research and Advanced Studies, ICREA

**Abstract**

It is a matter of consensus that the ability to efficiently use current and future High Performance Computing systems is crucial for science, however, the reality is that the performance currently achieved by most of the parallel scientific applications is far from desired. Despite inter-process communication has already been a matter of study in many different works, it is a fact that their recommendations are not taken into account in most of computational model development processes, at least in the case of Earth Science. This work presents a methodology that aims to help scientists working with computational models using inter-process communication to deal with the difficulties they face when trying to understand their applications behaviour. Following a series of steps that are presented here, both users and developers will learn how to identify performance issues by characterizing applications scalability, identifying which parts present a bad performance and understand the role that inter-process communication plays. In this work, the Nucleus for European Modelling of the Ocean (NEMO), the state-of-the-art European global ocean circulation model, will be used as an example of success. It is a community code widely used in Europe, to the extent that more than a hundred million core hours are used every year in experiments involving NEMO. In the analysis exercise it is shown how to answer the questions of where, why and what is degrading model's scalability, and how this information can help developers in finding solutions that will mitigate their eventual issues. This document also demonstrates how performance analysis carried out with small size experiments using limited resources can lead to optimizations that will impact bigger experiments running on thousands of

*Preprint submitted to Journal of Computational Science* *March 10, 2017*

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# ELPiN

- Removing the land-only processes in the smart way.

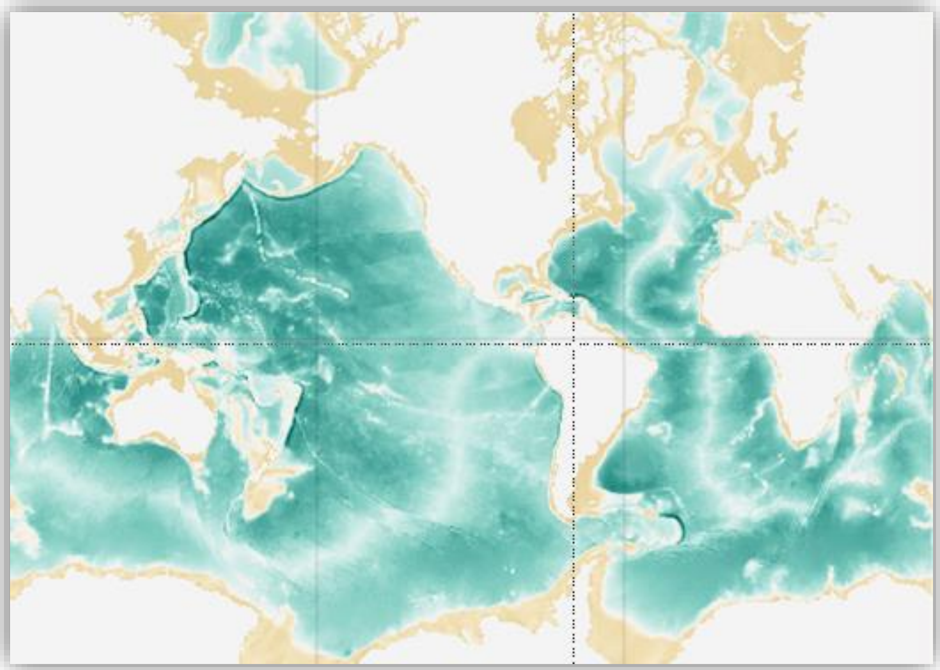

# ELPiN

- Removing the land-only processes in the smart way.

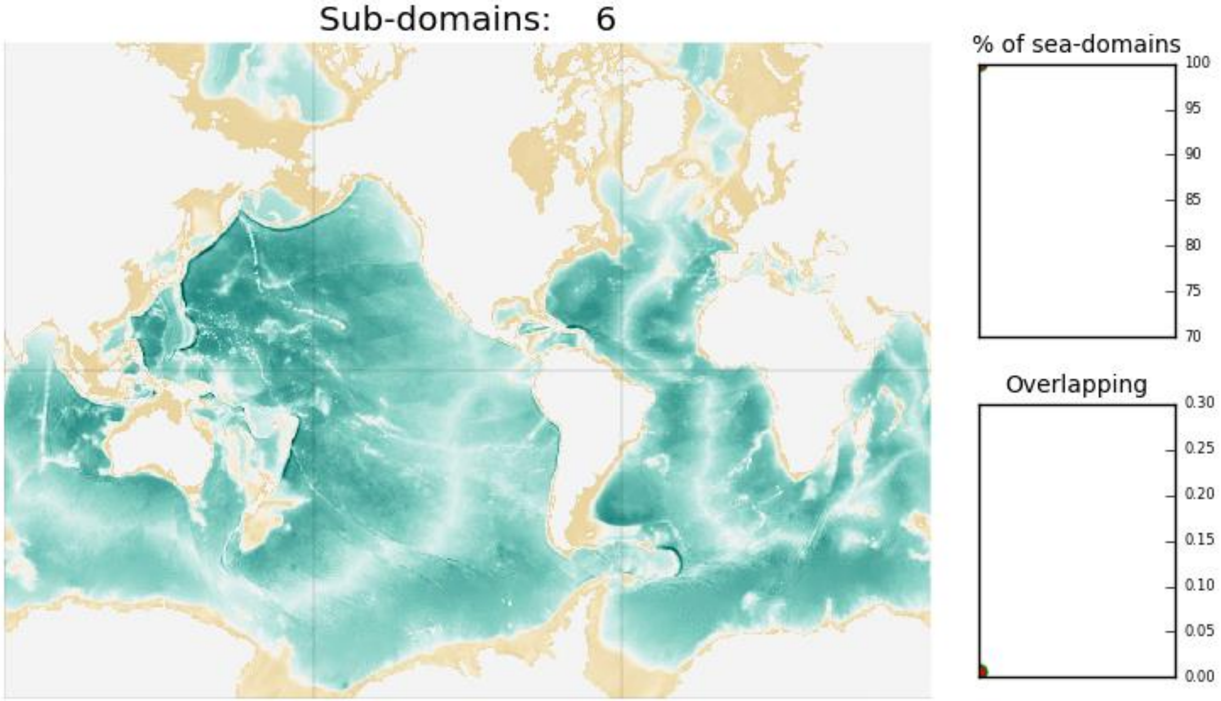


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
UAB Universitat Autònoma de Barcelona

# ELPiN

- Impact on ORCA025 simulations

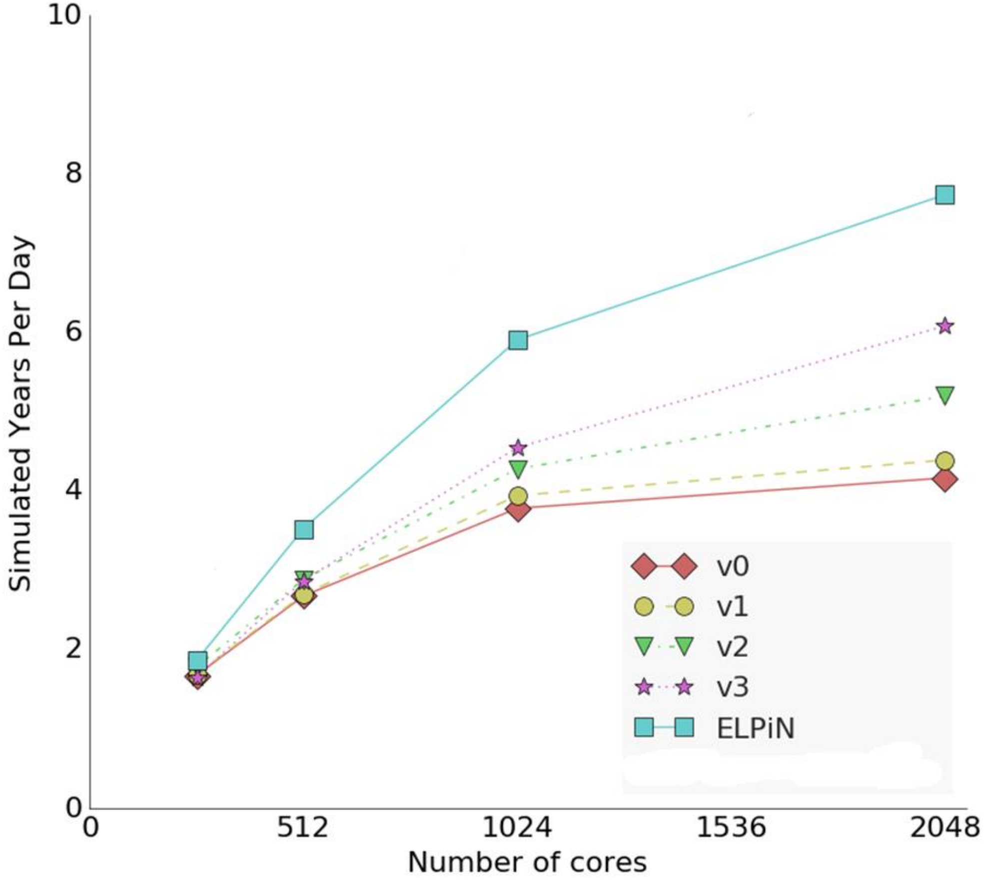

# Impact of our work

- This tool has been implemented in the EC-Earth **production** workflow. CMIP-6 simulations with EC-Earth will be using it and therefore saving millions of computing hours.

- Presented at **ICCS 2017** @ Zurich

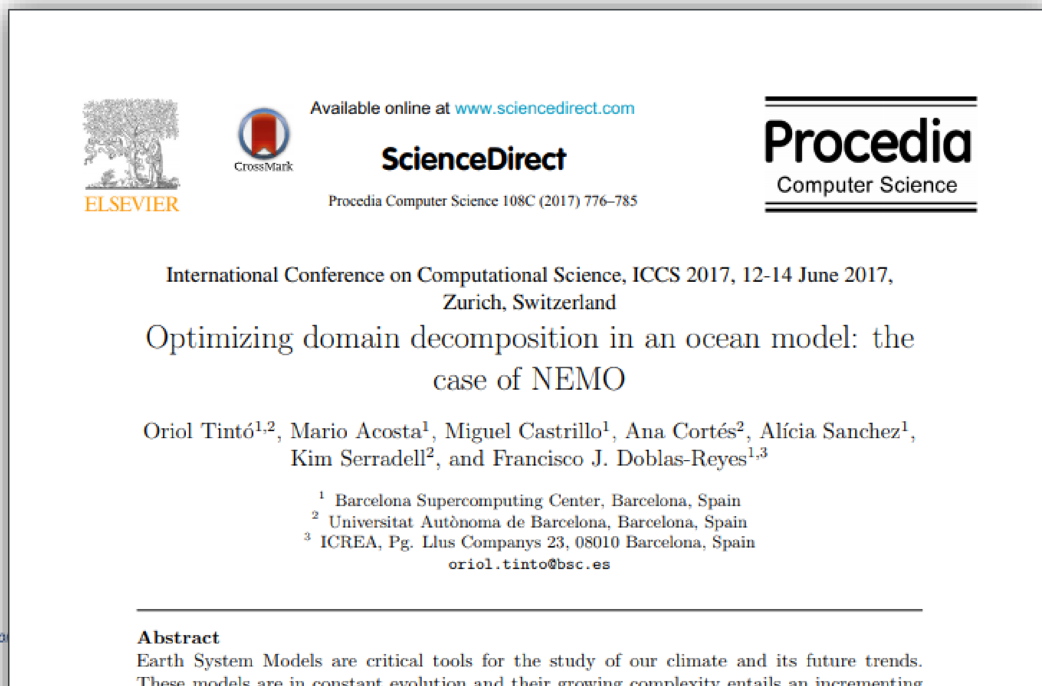

Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 776–785

Procedia Computer Science

ELSEVIER

CrossMark

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

Optimizing domain decomposition in an ocean model: the case of NEMO

Oriol Tintó[1,2], Mario Acosta[1], Miguel Castrillo[1], Ana Cortés[2], Alícia Sanchez[1], Kim Serradell[2], and Francisco J. Doblas-Reyes[1,3]

[1] Barcelona Supercomputing Center, Barcelona, Spain
[2] Universitat Autònoma de Barcelona, Barcelona, Spain
[3] ICREA, Pg. Llus Companys 23, 08010 Barcelona, Spain
oriol.tinto@bsc.es

**Abstract**
Earth System Models are critical tools for the study of our climate and its future trends.
These models are in constant evolution and their growing complexity entails an incrementing

# Exploring the use of mixed Precision

Exponent

$$1.2345 = \underbrace{12345}_{\text{Significand}} \times \underbrace{10}_{\text{Base}}^{\overbrace{-4}^{\text{Exponent}}}$$

Significand Base

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Exploring the use of mixed Precision

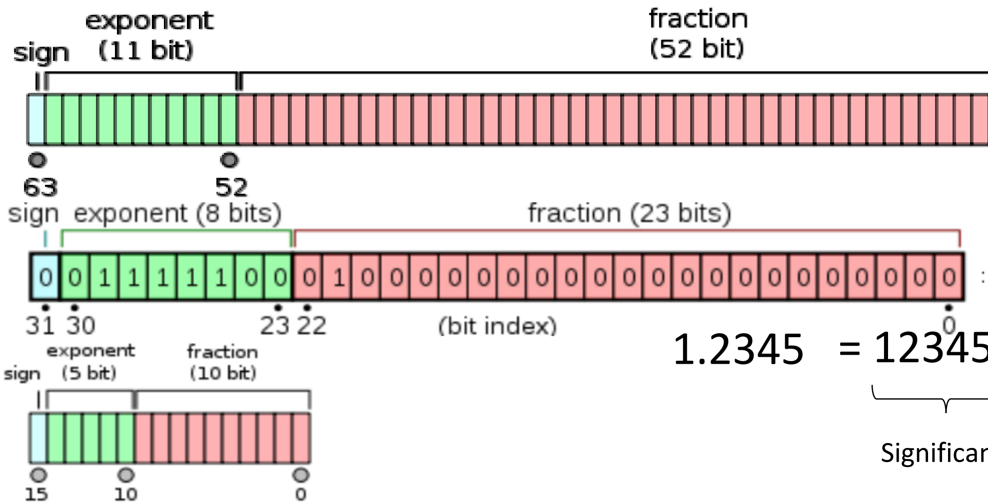


Exponent

$$1.2345 = 12345 \times 10^{-4}$$

Significand Base

# Exploring the use of mixed Precision

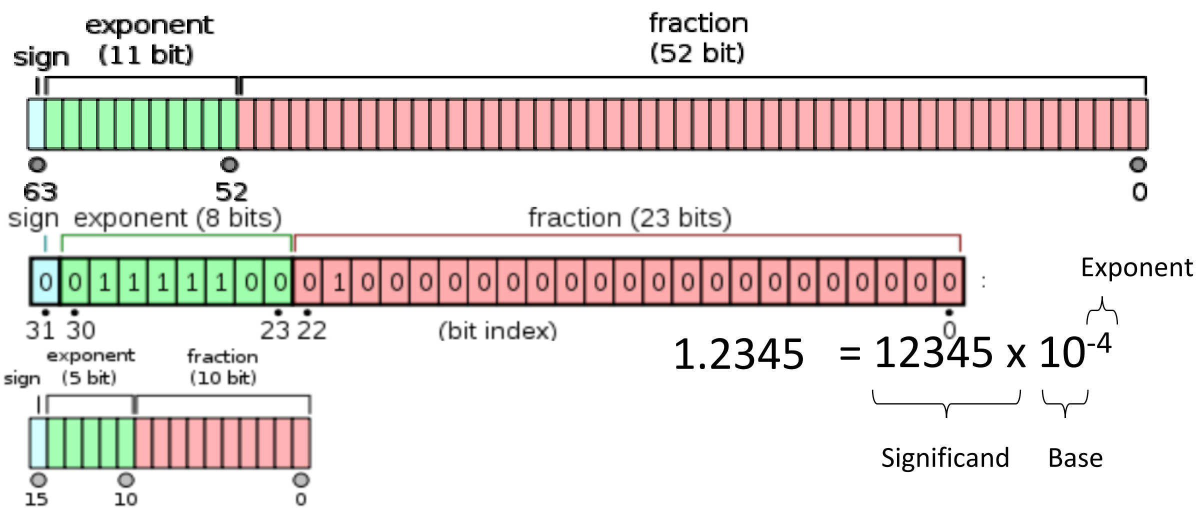


- Which precision is needed for the data that we want to represent?

# Exploring the use of mixed Precision

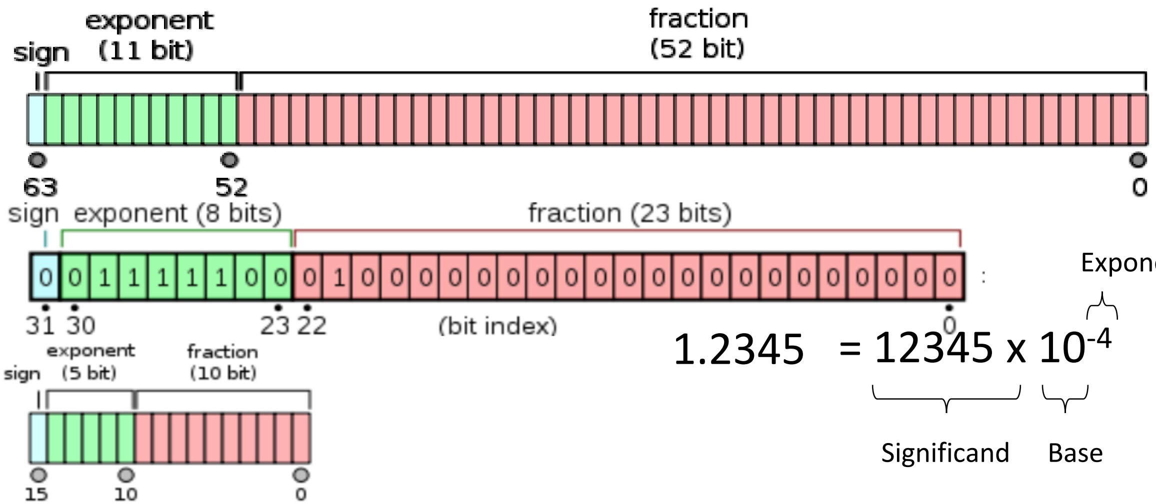


~$10^{-16}$

~$10^{-7}$

~$10^{-3}$

Exponent

$1.2345 = 12345 \times 10^{-4}$

Significand Base

- Which precision is needed for the data that we want to represent?

# Exploring the use of mixed Precision

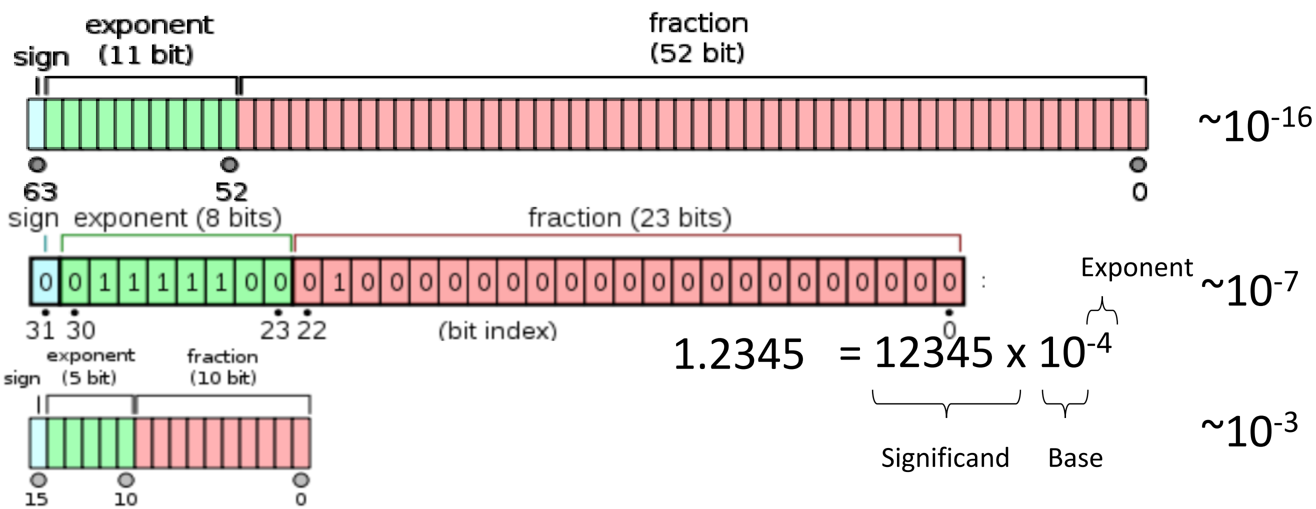


- Which precision is needed for the data that we want to represent?

# Exploring the use of mixed Precision

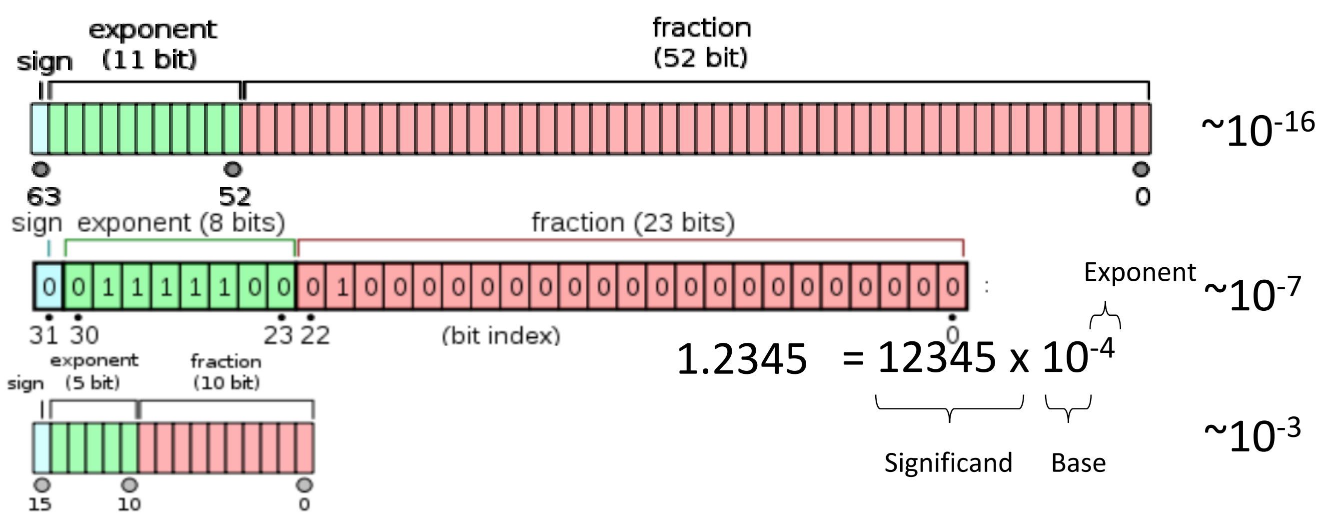


- Which precision is needed for the data that we want to represent?

# Exploring the use of mixed Precision

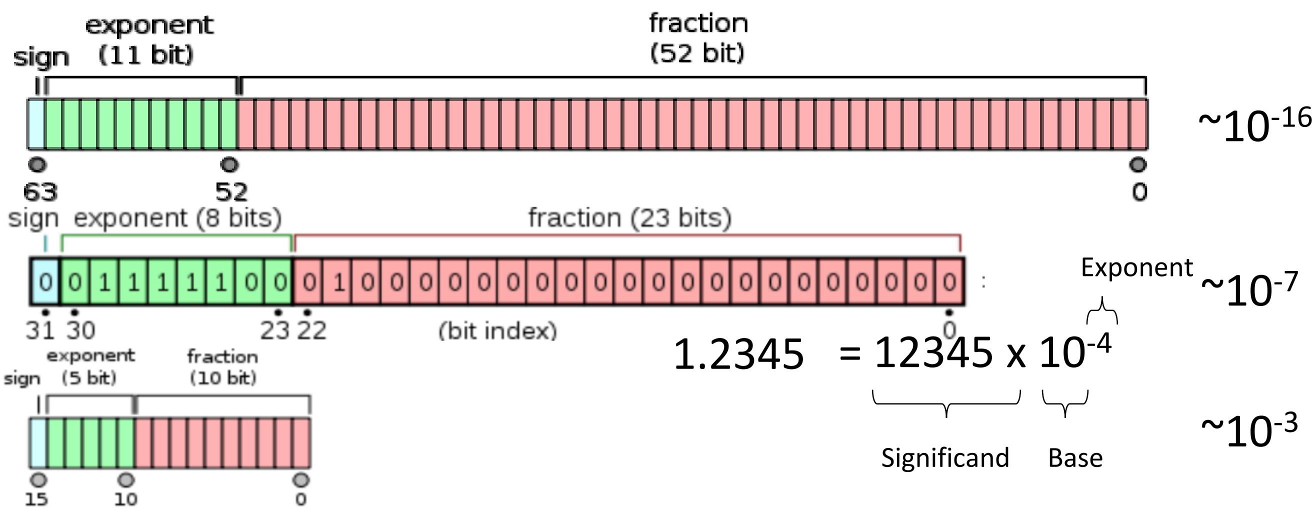


- Which precision is needed for the data that we want to represent?

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

# Exploring the use of mixed Precision

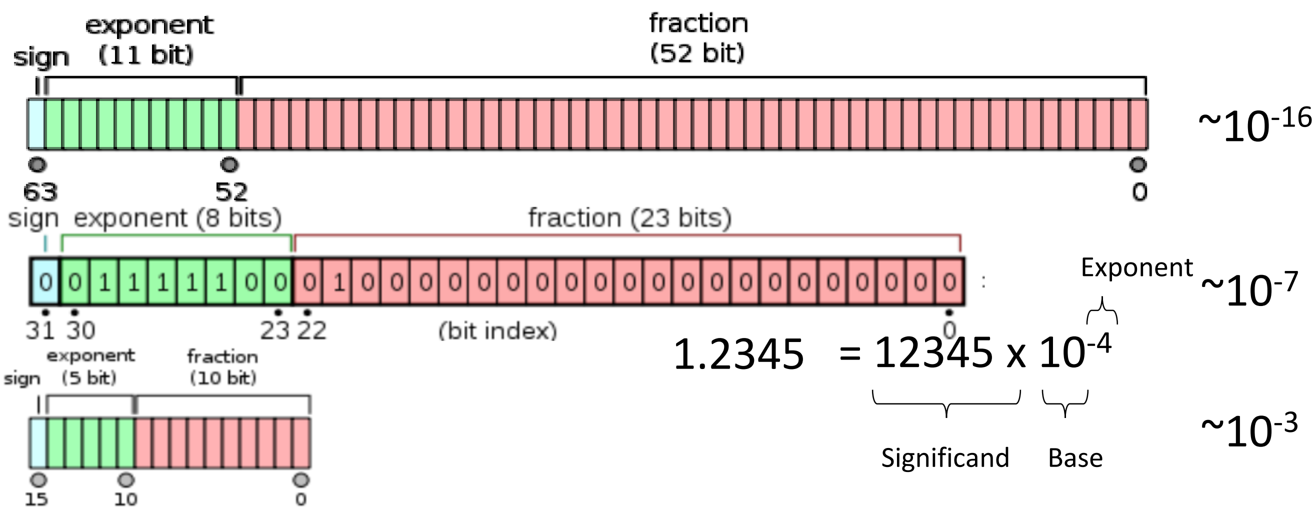


- Which precision is needed for the data that we want to represent?

Satellites can measure sea surface temperature with an uncertainty of 0.3 ºC and surface wind with an uncertainty of 1 m/s.

*- Remote Sensing of European Seas - V.Barale, M.Gad*

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
UAB Universitat Autònoma de Barcelona

# Exploring the use of mixed Precision

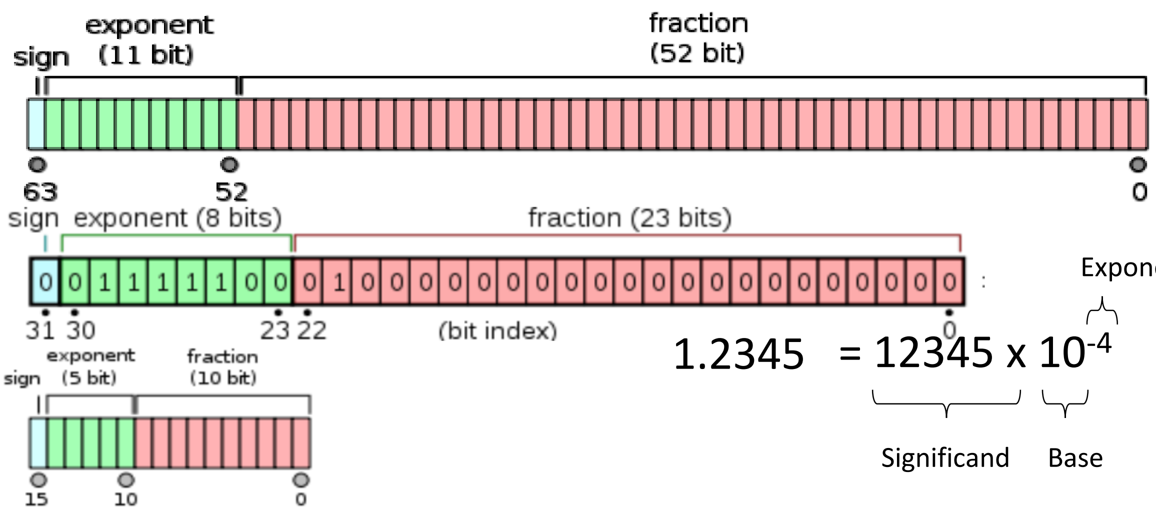


$\sim 10^{-16}$

$\sim 10^{-7}$

Exponent

$1.2345 = 12345 \times 10^{-4}$

Significand Base

$\sim 10^{-3}$

- Which precision is needed for the data that we want to represent?

Satellites can measure sea surface temperature with an uncertainty of 0.3 ºC and surface wind with an uncertainty of 1 m/s.

*- Remote Sensing of European Seas - V.Barale, M.Gad*

$\sim 10^{-1}$

# Exploring the use of mixed Precision

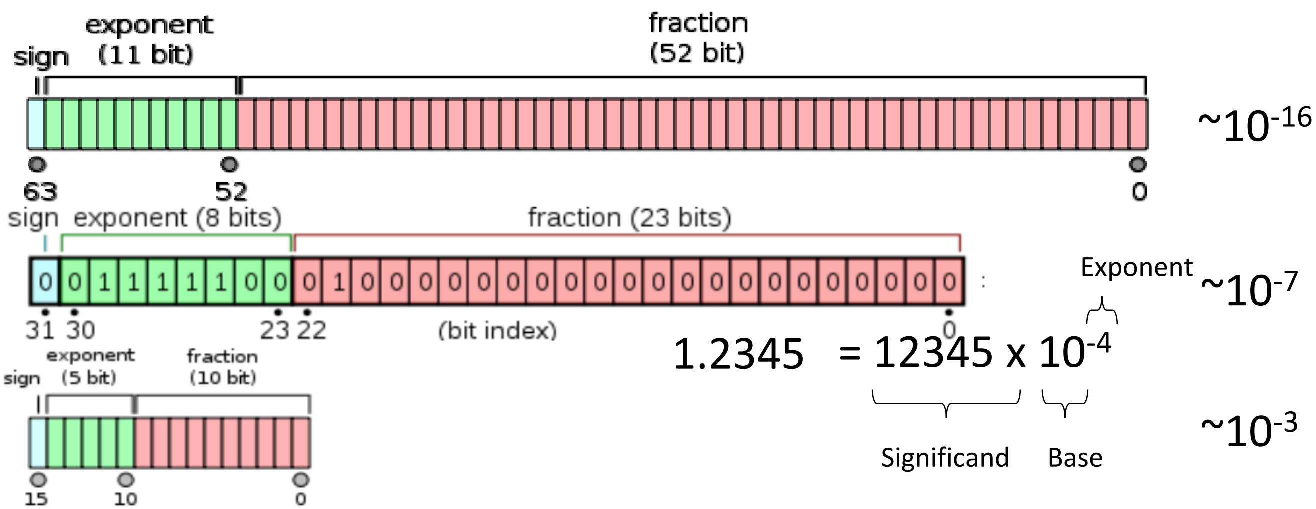


- Which precision is needed for the data that we want to represent?

Satellites can measure sea surface temperature with an uncertainty of 0.3 ºC and surface wind with an uncertainty of 1 m/s.
*- Remote Sensing of European Seas - V.Barale, M.Gad*

$\sim 10^{-1}$

To represent data with this level of uncertainty, using **half-precision (16-bit) should be enough**. Instead, double precision (64-bit) is the standard.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
Universitat Autònoma de Barcelona

# Reduced Precision Emulator

- Fortran Library developed by the Atmospheric, Oceanic & Planetary Physics group, within the University of Oxford Department of Physics

Geosci. Model Dev., 10, 2221–2230, 2017
https://doi.org/10.5194/gmd-10-2221-2017


Geoscientific Model Development
Open Access

## rpe v5: an emulator for reduced floating-point precision in large numerical simulations

**Andrew Dawson and Peter D. Düben**

Atmospheric, Oceanic & Planetary Physics, Department of Physics, University of Oxford, Oxford, UK

*Correspondence to:* Andrew Dawson (andrew.dawson@physics.ox.ac.uk)

Received: 20 September 2016 – Discussion started: 11 November 2016
Revised: 11 April 2017 – Accepted: 4 May 2017 – Published: 16 June 2017

**Abstract.** This paper describes the rpe (reduced-precision emulator) library which has the capability to emulate the use of arbitrary reduced floating-point precision within large numerical models written in Fortran. The rpe software allows model developers to test how reduced floating-point precision affects the result of their simulations without having to

in supercomputing power have not come from increasingly fast processors but from increasing the number of individual processors in a system. This has led to HPC applications being redesigned so that they scale well when run on many thousands of processors, and effort will be required to make sure models are efficient on the massively parallel su-

# Reduced Precision Emulator

- Fortran … y Physics group, v…

## Overview

The library contains a derived type: `rpe_var`. This type can be used in place of real-valued variables to perform calculations with floating-point numbers represented with a reduced number of bits in the floating-point significand.

## Basic use of the reduced-precision type

The `rpe_var` type is a simple container for a double precision floating point value. Using an `rpe_var` instance is as simple as declaring it and using it just as you would a real number:

```
TYPE(rpe_var) :: myvar

myvar = 12
myvar = myvar * 1.287    ! reduced-precision result is stored in `myvar`
```

## Controlling the precision

The precision used by reduced precision types can be controlled at two different levels. Each reduced precision variable has an `sbits` attribute which controls the number of explicit bits in its significand. This can be set independently for different variables, and comes into effect after it is explicitly set.

```
TYPE(rpe_var) :: myvar1
TYPE(rpe_var) :: myvar2

! Use 16 explicit bits in the significand of myvar1, but only 12 in the
! significand of myvar2.
myvar1%sbits = 16
myvar2%sbits = 12
```

Barcelona Supercomputing Center
Centro Nacional de Sup…

…utònoma …lona

# Method

# Method

**Implement the emulator**

- Replace REAL declarations
- Fix parameter declarations
- Solve calls to external routines
- Fix function calls with hardcoded reals
- Fix read and write statements
- Some more minor issues

**Design the tests**

- Define the error
- Minimize the number of tests
- Generate the namelists

**Execute tests**

- Manage all the different tests (eventually thousands)
- Handle dependencies
- Collect results

**Assimilate results**

- Convert the results to something meaningful and useful

# Implementing the emulator

# Implementing the emulator

- After some months lost manually implementing the emulator…
  a Python tool to automatize the process was created.

# Implementing the emulator

- After some months lost manually implementing the emulator…
  a Python tool to automatize the process was created.

Oriol Tintó-Prims / implement-rpe-to-NEMO

This project Search

README.md

## RPE implementation tool

This tool is intended to automatize the processes of implementing the Reduced Precision Emulator to a computational model.

Even it was first designed to be used for the **NEMO** model, it has evolved to be useful for **any computational model written in FORTRAN**.

The script automatically:

- Replace REAL variables.
- Track dependencies between variables.
- Fix parameter declarations.
- Fix routine/function calls.
- Fix WRITE and READ statements.
- Add module to read precisions from namelist.
- Add precision assignment for each variable.

It can be also used to:

- Generate namelists for different cases.

# Preparing tests to analyze the impact of the numerical precision

# Preparing tests to analyze the impact of the numerical precision

- Models have thousands of variables, and individual screening of each variables may be expensive and unnecessary.
- Minimizing the number of tests to do:

$$\boldsymbol{P}(n) = \sum_{x=0}^{n-1} (1-p)^x \cdot p$$

$$\boldsymbol{N^{o}Tests} = \frac{N}{n} + \frac{N}{n} \cdot P(n) \cdot n = N \cdot \left(\frac{1}{n} + P(n)\right)$$

Where:

- **p**: Probability of a variable to need more than Z bits of precision
- **P**: Probability of having a variable needing more than Z bits of precision in a test with **n** members
- **n**: Size of the groups
- **N**: Total number of variables

# Preparing tests to analyze the impact of the numerical precision

- Models have thousands of variables, and individual screening of each variables m
- M

Whe
- p:
- P:
- n:
- N

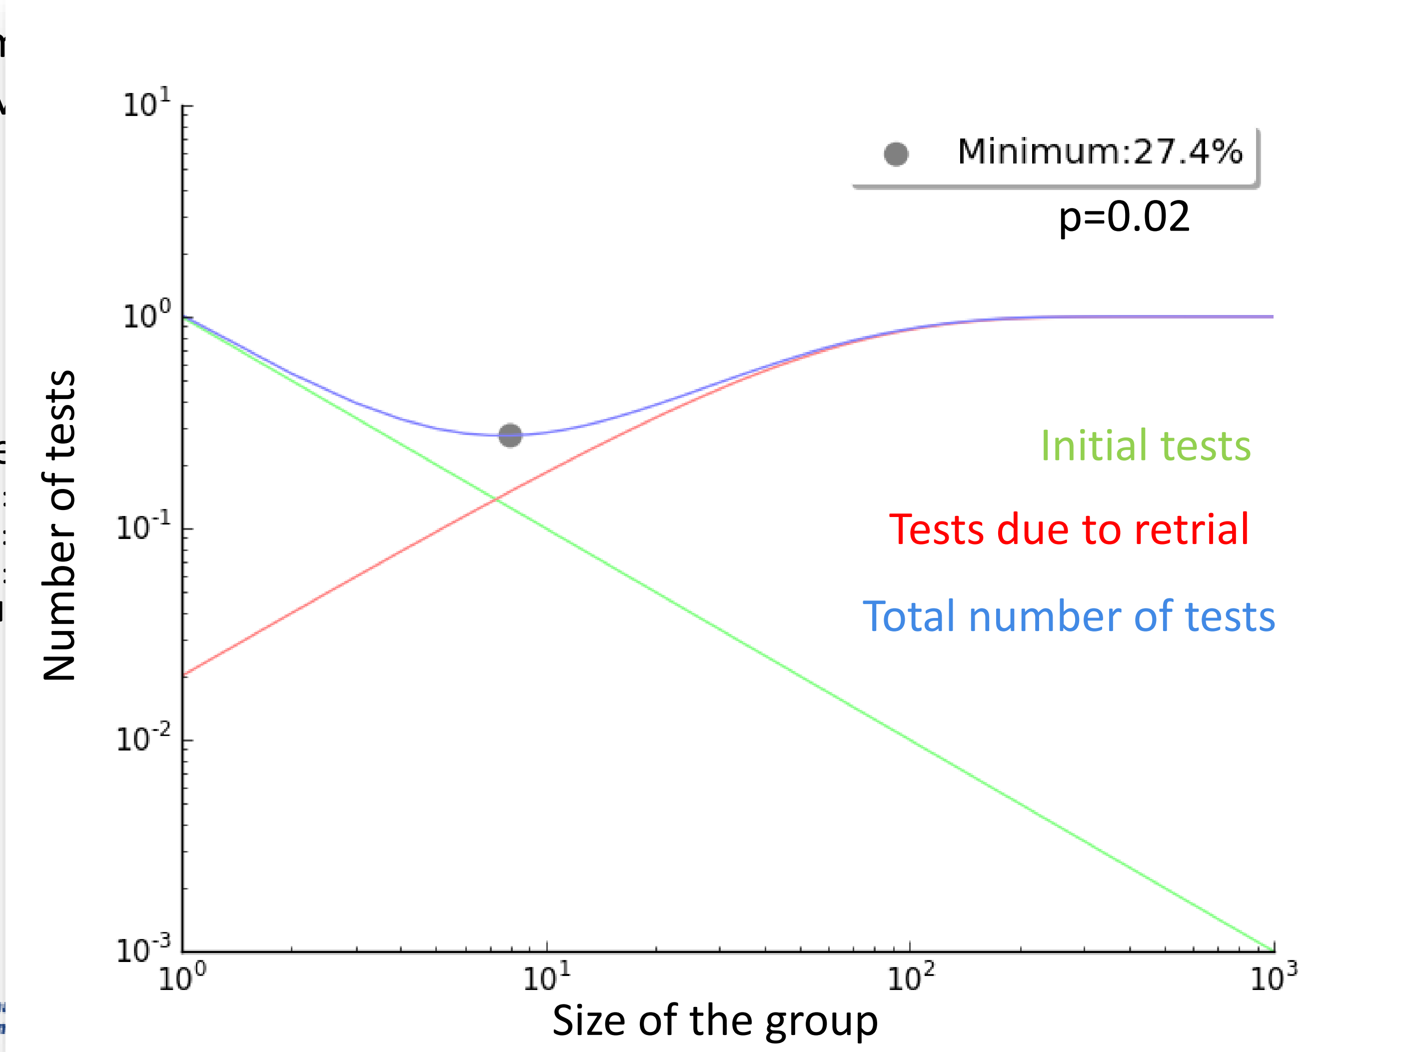


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

de Barcelona

# Executing the tests

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Executing the tests

- A case test with a **GYRE** configuration has **3509** variables to be tested.
- With **autosubmit** we designed the following workflow:

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

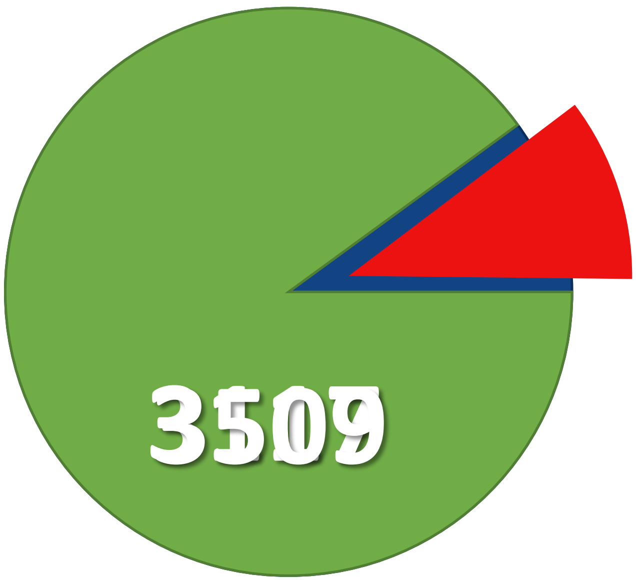

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

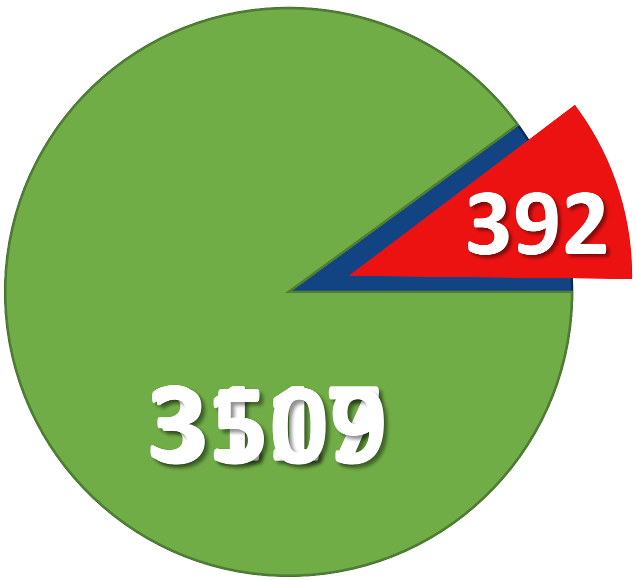


**89%** of the variables **can use 10 bit** for the significand without adding significant error

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos UAB Universitat Autònoma de Barcelona

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

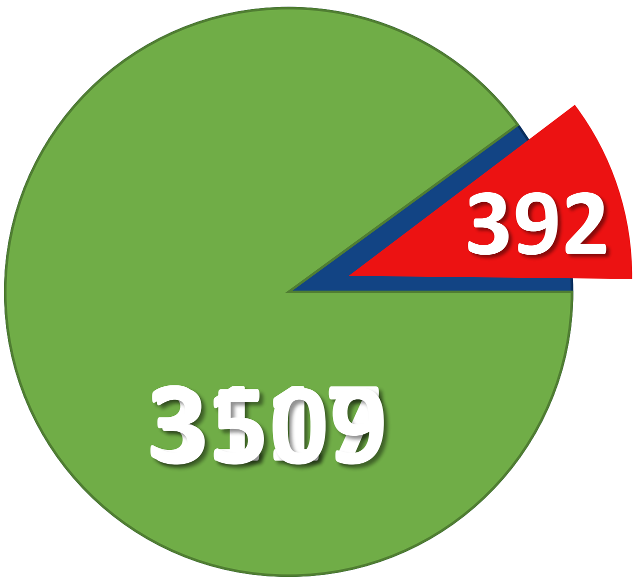


**89%** of the variables **can use 10 bit** for the significand without adding significant error

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

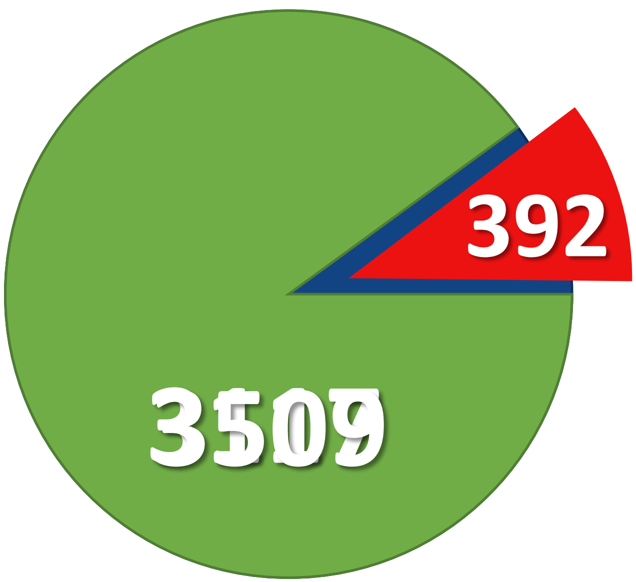


**89%** of the variables **can use 10 bit** for the significand without adding significant error

# Assimilate the results

- After all the tests are completed, we obtain a list of the effect that reducing the precision of each variable produces to the outputs.

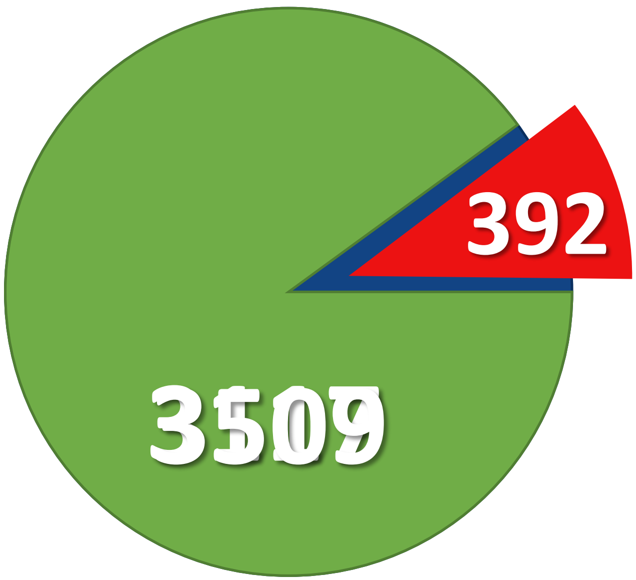


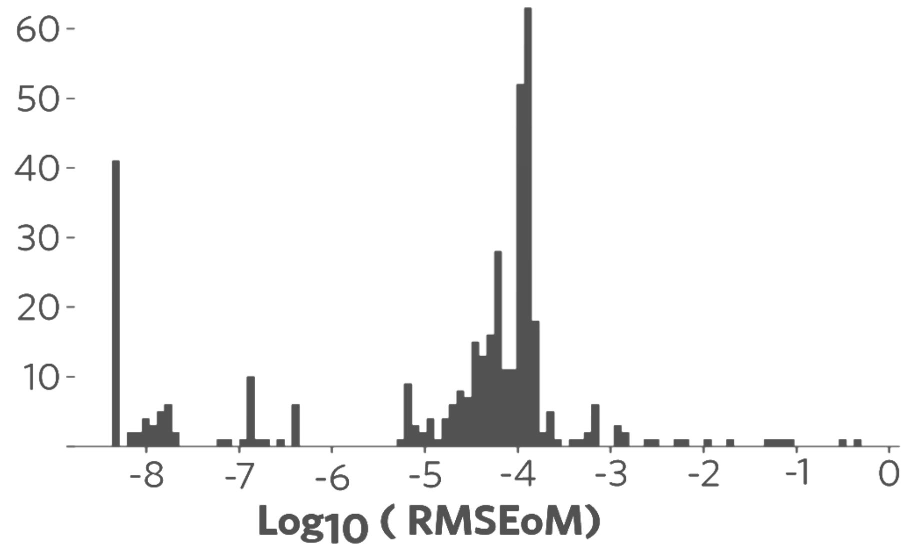


**89%** of the variables **can use 10 bit** for the significand without adding significant error

# Exploring the use of mixed Precision

- Potential Impact?

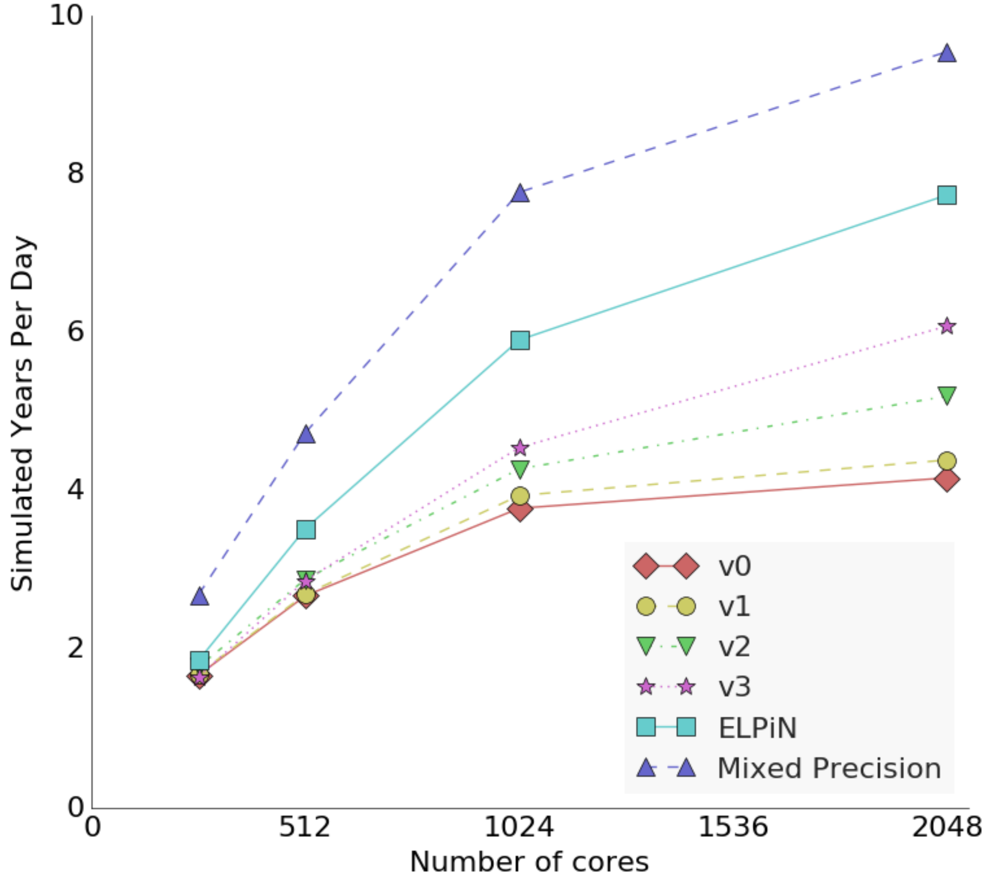


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

# Collaborations

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

# Collaborations

- Three months international secondment at **University of California Santa Cruz**.
- **NEMO HPC** and development teams.
- **EC-Earth** community.
- University of Oxford's Department of Physics , group of **Predictability of Weather and Climate**
- RIKEN/AICS

# Conclusions

Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

# Conclusions

- The main research objectives of the thesis have been fulfilled, even surpassing the initial scope and being useful for other Earth System models.
- The main objective of the last year of the thesis will be the publication of at least two articles. Additionally, an actual mixed-precision implementation of the NEMO model will be done.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

caos

UAB Universitat Autònoma de Barcelona

# Thank you

oriol.tinto@bsc.es